

Towards an Account of Basic Final Value
Timothy Perrine
Inquiry, DOI: 10.1080/0020174X.2023.2262758

Abstract: Ordinary and philosophical thought suggests recognizing a distinction between two ways something can be of final value. Something can be of final value in virtue of its connection to other things of value (“non-basic final value”) or something can be of final value regardless of its connection to other things of value (“basic final value”). The primary aim of this paper is to provide an account of this distinction. I argue that we have reason to draw this distinction as it helps avoid certain problems. I criticize accounts of this distinction due to Warren Quinn, Fred Feldman, and Michael Zimmerman. I then provide my own positive account which incorporates several of the insights of those accounts while avoiding their pitfalls. I conclude by relating my account to issues concerning partiality and appropriate attitudes.

Some things are valuable for their own sake—being happy or experiencing pleasure perhaps. Other things are valuable because their consequences are things of value—having money being the standard example. This distinction in value is the distinction between what is sometimes called *final* or *intrinsic* value and *instrumental* value. This distinction is ancient (cf., e.g., Plato’s *Republic* II, 357a-358b, Aristotle’s *Nicomachean Ethics*, I.6 1096b, Cicero, *OME*, Book I, 29, 42) and can frequently be found in ordinary thinking.

It is natural to recognize a further distinction between things of final value. Some things are of final value, but only because of their connection to other things of final value. For instance, many people think that visiting relatives is of final value. But insofar as it is, it is only because of its connection to other specific things of final value, such as the particular interactions one has with one’s relatives during the visit. By contrast, other things are of final value but not because of their connections to other things of final value. Perhaps my enjoyment of a beer at a particular time is of final value. But if so, it is not because of its connection to other things of final value.¹

To introduce some standard terminology, when something is of final value because of its connection to other things of final value, it is of “non-basic” final value; by contrast, when something is of final value regardless of its connections to other things of final value, it is of “basic” final value. Not only is it natural to recognize this distinction, as several philosophers have argued (e.g. Harman (1967), Feldman (2000), Zimmerman (2001)), it is *important* to recognize it as well, since recognizing it helps avoid certain problems.

But recognizing a distinction between basic and non-basic final value falls short of providing an account of that distinction. The overall aim of this paper is to provide an account of the distinction between basic and non-basic final value. Taking non-basic final value as final value that is not basic, giving an account of that distinction amounts to giving an account of basic final value; thus, I will give an account of basic final value. I begin, in section I, by laying out some preliminary assumptions about final value. In section II, I cover various reasons why it is important to draw a distinction between basic and non-basic final value to avoid certain problems. In section III, I argue against accounts of this distinction due to Warren Quinn, Fred

¹ As the examples illustrate, by things of value, I have in mind facts, events, states of affairs, or ways the world could be. I don’t have in mind specific concrete objects, like a frying pan. See fn. 3 for further relevant discussion.

Feldman, and Michael Zimmerman. In section IV, I offer my own account. I show how it avoids the problems mentioned in section II and avoids the pitfalls of alternative accounts articulated in section III. Finally, in section V, I explore how my views about value relate to valuing things for their own sake as well as issues of partiality.

I. Preliminaries

The kind of value I am interested in is that kind of value that is valuable for its own sake or in and of itself. Different authors use different terms to refer to this kind of value, including ‘intrinsic value,’ ‘intrinsic goodness,’ ‘final value,’ and ‘final goodness.’ I will use the term ‘final value.’ To avoid confusion when discussing other authors, I will enforce terminological consistency by using brackets, [], to indicate where I’ve modified their text. I do not assume that when something is of final value, it has final value solely in virtue of its intrinsic properties. And, of course, I assume that there is such a kind of value.

I assume that there is *some* sort of substantive normative connection between final value and pro-attitudes (and final disvalue and con-attitudes). By pro-attitudes, I mean favorable attitudes directed at things. Some examples include: valuing ϕ , favoring ϕ , caring about ϕ , desiring ϕ , respecting ϕ , being pleased about ϕ , etc. (For longer lists see Adams (2007: 17-9) and Nozick (1981: 429ff.)) The idea that there is *some* sort of normative connection is widely embraced by philosophers of otherwise different perspectives and views.² For purposes here, I assume that the best way to understand this connection is in terms of the “appropriateness” or “fittingness” of pro-attitudes. That is, I’ll be assuming that *something* like the following principles are correct:

Final Value: ϕ is of final value if and only if it is appropriate for agents to bear pro-attitudes towards ϕ for its own sake.

Final Disvalue: ϕ is of final disvalue if and only if it is appropriate for an agent to bear con-attitudes toward ϕ for its own sake.

There are other ways of spelling out the connection between final value and pro-attitudes. For instance, other authors might speak of when certain pro-attitudes are required (e.g., Zimmerman (2001), Lemos (2011)) or what reasons agents have for having certain pro-attitudes (e.g., Scanlon (1998)). As I’ll briefly discuss in section V, there are reasons for preferring spelling out the connection between final value and pro-attitudes in terms of appropriateness or fittingness instead of these other ways.

I also assume that the primary bearers of final value and disvalue are states of affairs. For simplicity, I will identify states of affairs with propositions (following Chisholm (1976: chp. 4)). I assume the following about propositions, and thus states of affairs. First, propositions are *abstracta*. Second, propositions are closed under Boolean operators—for any proposition, there is the proposition that negates it; for any two propositions, there is a proposition that is their conjunction. I’ll use the phrase ‘complex states of affairs’ to refer to states of affairs that are conjunctions of states of affairs. I’ll use the phrase ‘states of affairs’ to be neutral as to whether I’m talking about a complex state of affairs or a non-complex one. Third, many propositions are attributions of properties or relations to particular objects or things; in this sense, they are complex entities. Propositions may be partially distinguished on the basis of which properties or relations they attribute to which objects or things.

² See, for instance, Brentano (1902: 15f.), Ross (1939: 279, 282), Ewing (1947: 146ff.) Chisholm (1986: 47ff.), Anderson (1993: 2-3), Lemos (1994: 6ff.), Scanlon (1998: 78ff.), Zimmerman (2001), Audi (2004: 125). Merely maintaining that there is a substantive normative connection in no way identifies which property cluster is more fundamental. Further, this assumption does not require the particular kind of buck-passing popularized by Scanlon that is open to the wrong-kind of reason objection. For discussion of this last point, see Zimmerman (2010).

Some propositions have a fairly-fine grain. Even if several propositions are, in some crude or loose sense, about the same event they may still fail to be identical. By way of illustration, consider the following sequence:

- (1) Ava shrieked.
- (2) Ava shrieked loudly.
- (3) Ava shrieked loudly with an average decibel of D.

Even if (1)-(3) are about the same event—in some broad sense of ‘event’—I assume that they are nonetheless distinct propositions.

Some authors have suggested that individual objects—like, perhaps, a mink coat (Korsgaard (1983)), or Abraham Lincoln’s pen (Kagan (1998))—can be bearers of final value. That position is plausible. But for reasons others have discussed (e.g., Zimmerman (2001: 33-9)), it is plausible that when particular things are of final value, it is because of their connection to states of affairs that are of final value. No individual object is of final value without there being a state of affairs about that object that is of final value. Thus, in what follows, I’ll focus on states of affairs as I understand them, not individual objects, as the most promising candidate for being the primary bearers of value. (For a version of the opposite view, see Rabinowicz and Rønnow-Rasmussen (2000, 2003).)³

II. Reasons to Draw the Distinction

A preliminary way of drawing the distinction between basic and non-basic final value is this. Some states of affairs are of final value but only because of their relation to other states of affairs of final value. (For instance, perhaps a state of affairs might be of final value because it has, as a part, another state of affairs of value.) Other states of affairs are of final value but *not* because of their relation to other states of affairs of final value. The former are of “non-basic” final value; the latter are of “basic” final value.

This preliminary way of drawing the distinction falls short of being an adequate account of the distinction. Compare: an innate trait is a trait something has in virtue of its ancestry, whereas an acquired trait is a trait something has in virtue of its environment. This preliminary way of drawing the distinction falls short of being an adequate account of the distinction between innate traits and acquired traits; we might seek a fuller account of the distinction. So too, even with this preliminary way of drawing the distinction in hand, we might seek a fuller account of the distinction between basic and non-basic final value. Further, different accounts of this distinction will offer different refinements or explications of this preliminary way of marking the distinction. But there are several reasons for thinking that an adequate axiology *should* mark this distinction. The aim of this section is to articulate some of those reasons.⁴

³ I focus on relationships between states of affairs of final value. Some authors might be interested in the relationships between concrete things that are of final value. For instance, Langton (2007) suggests that something—e.g. a wedding ring—might be of final value but in virtue of something else—e.g., a personal relationship (see also Korsgaard (1983)). These cases are sometimes called cases of “extrinsic final value,” since the wedding ring is of final value, but the “source” of its value is extrinsic to it. Since these cases involve individual objects, they are not the focus of my discussion. It’s an interesting question how exactly views about extrinsic final value interact with my proposals here. I don’t see any grand conflict. After all, if an object is of extrinsic final value, there’s a state of affairs about that object and its source of value that is, presumably, of final value. But I hope to explore this question in future work. For a negative view about how they might interact, see Bradley (2006); for a critical discussion of extrinsic final value in general, see Tucker (2016).

⁴ The reasons are drawn from Harman (1967), Feldman (2000), and Zimmerman (2001), though the presentation is mine. See also Moore (1912: 36ff.) and Ross (1930: 68ff.).

However, one more clarificatory comment is in order. The distinction between basic/non-basic final value is both similar and dissimilar to the distinction between instrumental/final value. When something is of instrumental value, there is a “connection” between it and other things of value or disvalue. Further, an explanation of the value it has will appeal to the value (or disvalue) of those things. Similarly, when something is of non-basic final value, there is a “connection” between it and other things of final value. Further, an explanation of its value will appeal to the value of those other things. But there are some important differences between instrumental value and non-basic final value. First, when something is of instrumental value, the “connection” it has to other things of value is having, as a causal consequence, that things of value (disvalue) obtain (do not obtain). That need not be the case for non-basic final value. In fact, I’ll suggest below that the relevant connection between non-basic final value and final basic value should be understood mereologically—in terms of parts—and not necessarily in terms of causal consequence. Second, when something is of non-basic final value it follows that it is of final value. Thus, given the principle *Final Value*, when something is of non-basic final value, it follows that it would be appropriate to adopt certain pro-attitudes towards it for its own sake. By contrast, when something is of instrumental value, it does not follow that it is of final value—it may be, it may not be.

A. *Ontological Flexibility*

There are various theories of value, different axiologies. It is natural to seek a formulation of a theory of value—a formulation that states clearly what is of value or not, given that theory. Such a formulation should do two things. First, it should be suitable for philosophical and ordinary purposes. Second, it should be true to the theory that it is meant to represent—neither being too restrictive nor too loose by the lights of that theory. However, if we neglect the distinction between basic/non-basic final value, there are some natural ways of formulating a theory of value that will not do those things.

I’ll illustrate these points with a simple form of hedonism. A stock formulation of hedonism—the kind you might find in an introductory text—is that only pleasure is of final value and only displeasure is of final disvalue. Filtering that statement through the idea that states of affairs are the primary bearers of value, we might formulate this simple form of hedonism as:

Simple Hedonism: A state of affairs is of final value only if it attributes pleasure (and nothing else) to something; a particular state of affairs is of final disvalue only if it attributes displeasure (and nothing else) to something.

Simple Hedonism has an important consequence. If there is a complex state of affairs that attributes more than just pleasure or displeasure to things, then *Simple Hedonism* implies that it is neither of any final value or disvalue.

Because of that consequence, *Simple Hedonism* is an inadequate formulation of this simple form of hedonism. First, it is inadequate for philosophical purposes. Philosophers are interested in evaluating, among other things, the final value of entire possible worlds, or outcomes of actions, or states of affairs representing a person’s life. Since those complex states of affairs attribute more than just pleasure or displeasure to things, *Simple Hedonism* implies that they are neither of any final value or disvalue. Second, it is inadequate for ordinary purposes. To use the earlier example, many think the state of affairs consisting of visiting one’s family can be of final value. But such a complex state of affairs will attribute more than just pleasure or displeasure to things. Thus, *Simple Hedonism* wrongly implies that such a complex state of affairs is not of any final value. Finally, it is too restrictive even by the lights of proponents of this simple form of

hedonism. For instance, consider a complex state of affairs that represents an evening with friends that is overall pleasant. Because the evening was overall pleasant, an adequate formulation of hedonism should imply that the complex state of affairs is of final value. But *Simple Hedonism* will not because the complex state of affairs will contain states of affairs that do not attribute just pleasure or displeasure. (E.g., it will include that Alice was drinking wine, or that euchre was played, etc.)

The basic problem here is one of “ontological flexibility.” An adequate formulation of a theory of value should allow a range of complex states of affairs to be of final value or disvalue. Formulations of axiologies like *Simple Hedonism* lack the required flexibility.

The distinction between basic and non-basic final value helps avoid this problem. An axiology like hedonism can be formulated first and foremost in terms of which states of affairs are of *basic* final value. Other states of affairs—like those representing a pleasant evening with friends, going to a musical concert, possible outcomes of actions, or even entire possible worlds—can then be of *non-basic* final value in virtue of their relation to states of affairs of basic final value. Drawing a distinction between basic and non-basic final value gives an axiology the ontological flexibility that is needed for philosophical and ordinary purposes.

B. Computing Final Value: Simple Over Counting

A lesson from the previous section is that axiologies should be formulated in such a way so that complex states of affairs can be of final value. But if complex states of affairs can be of final value, we would ideally know how the value of their parts contribute to the value of the whole.⁵ Further, if we do not draw the distinction between basic and non-basic final value, it may be easy to over count the value of certain states of affairs when determining the value of a whole. In fact, there are several ways we might over count value.

First, we might over count value when the value of a single part is counted more than once. By way of illustration, consider the following view: the value of a complex state of affairs is the sum of the values of all of its proper parts. This view has a well-known problem (cf. Zimmerman (2001: 154-5)). Suppose A, B, and C are states of affairs, each is a proper part of W (e.g. $W = A \& B \& C$), and A is of some value—say, 10 units—and B and C are not of any value, positive or negative. The proper parts of W include A, B, C, A&B, A&C, B&C. Further, intuitively, A&B and A&C are of 10 units of final value, given that A is of 10 units of value and neither B nor C are of any value or disvalue. But on this view of the value of complex states of affairs, the value of W will be 30 units—intuitively too high. Clearly, the reason why 30 units is too high is that the value of A was counted three times instead of just once.

There is a second, subtler way we might over count the value of things (cf. Feldman (2000: 322-3)). Several states of affairs might be about the same particular event but nonetheless non-identical. By way of illustration, suppose after a hard day at work George sits down and enjoys his favorite beer. All of the following states of affairs would then obtain:

- (4) At time interval *I*, George experiences some pleasure.
- (5) At time interval *I*, George experiences some pleasure in his beer.
- (6) At time interval *I*, George experiences a great deal of pleasure in his beer.

None of these states of affairs are identical. For instance, (6) entails (5), but not conversely; likewise, (5) entails (4) but not conversely. But, intuitively, each of these states of affairs are

⁵ I don't assume that the value of complex states of affairs is *always* determined exclusively by the value of its parts. Nonetheless, I *do* assume that sometimes the value of complex states of affairs is determined at least partially by the value of its parts. Further, I do not assume that if the value of a complex state of affairs is determined by the value of its parts, then it determined in a simple, additive way.

about the same event; it is merely that (6) contains more information than either (4) or (5). Even if (4)-(6) are all of final value, the final value of each should not be included in a computation of a complex state of affairs that includes each of them. For (4)-(6) are not about three *distinct* cases of pleasure; thus counting the value of each for a complex state of affairs would be over counting.

The distinction between basic and non-basic value helps avoid these types of problems of over counting. Some states of affairs are of non-basic value. But their non-basic value is *wholly* explained by certain states of affairs of basic final value. Thus, we would double count were we to count *both* the basic value of states of affairs and the non-basic value of other states of affairs whose value is wholly explained by the basic value of those states of affairs.⁶

C. *Identifying What is of Instrumental Value*

When a state of affairs is of instrumental value it is because of its causal consequences—either in bringing about states of affairs of value (disvalue) or preventing things of value (disvalue). But this raises an important question: *which* consequences of a state of affairs determine its instrumental value? If we fail to draw the distinction between basic and non-basic final value, we might be led to answer this question in a way that is unduly problematic.

The simplest way of determining instrumental value is this:

Simple Instrumental Value: The instrumental value of a state of affairs is determined by *all* of its consequences that are of final value or disvalue.

However, this simple view faces a problem of over counting similar to the problem in the previous section (see Harman (1967: 798-9), Gibbard (1973) for more formal presentations of the problem). Suppose the following state of affairs obtains:

(7) At time *t*, George drinks his beer.

Several states of affairs are a causal consequence of (7) including:

(4) At time interval *I*, George experiences some pleasure.

(5) At time interval *I*, George experiences some pleasure in his beer.

(6) At time interval *I*, George experiences a great deal of pleasure in his beer.

Plausibly, each of (4)-(6) is of final value. Thus, given *Simple Instrumental Value*, the instrumental value of (7) is partly determined by all of (4)-(6). But that's a mistake. To give each of those states of affairs a role in determining instrumental value of (7) would be to over count.

The problem here is similar to the last problem: in computing the value of something we want to avoid over counting. The difference is that in the previous section we were concerned with the final value of some complex state of affairs; here we are concerned with the instrumental value of some state of affairs. Despite this difference, the solutions are similar. Some consequences are of non-basic value and their value is wholly explained by other consequences of basic value. It would be a mistake to count the value of both.

D. *Constraint*

The idea that some states of affairs have final value in a “basic way” that explains the final value of other states of affairs helps resolve some of these problems (or, more modestly, are the first steps of some promising ways of resolving these problems). But ideally we would have an

⁶ For my discussion here, I want to leave open the possibility that there are some states of affairs that are not of basic value but whose value is not fully explained by states of affairs of basic value. If there are such states of affairs, then a full reckoning of the final value of a complex state of affairs—like an entire possible world—might have to take them into consideration as well. For an opposite approach, see Feldman (2000: 333). (Thanks to a reviewer and Peter Finocchiaro for discussion of this topic.)

account of this distinction. In giving an account of this distinction, we should follow the following constraint.

The distinction between basic and non-basic value often concerns explanations between the value of states of affairs. And the above problems illustrate that we already have *some* intuitions about basic and non-basic value. These intuitions are mostly comparative and conditional. That is, they are intuitions about how if one state of affairs is of basic value, then others are (at best) not of basic value.

Here are some examples. Suppose that a state of affairs *S* is of final value while *S** is neutral. Then, intuitively, the state of affairs of *S&S** is not of *basic* final value. Likewise, consider the following states of affairs:

- (4) At time interval *I*, George experiences some pleasure.
- (5) At time interval *I*, George experiences some pleasure in his beer.
- (6) At time interval *I*, George experiences a great deal of pleasure in his beer.

Suppose that (6) is of basic final value. Then, intuitively, neither (4) nor (5) would be of basic final value as well. To be clear, the intuition here is not that (6) *is* of basic final value or not. (Maybe that is intuitive; maybe it is not.) Rather, what is intuitive is a conditional: were (6) of basic final value, then neither (4) nor (5) would be.

An analogy may be useful here. Suppose Liu prefers dumplings to noodles, and prefers noodles to hotpot. But, it turns out, Liu also prefers hotpot to dumplings. According to some decision theorists, it is intuitive that Liu's preferences are irrational. Their intuition is not that Liu's *specific* preference for (e.g.) dumplings over noodles is irrational. Rather, the intuition is that if certain of Liu's preferences are rational, then others are not. More metaphorically, the rationality of some preferences "constrains" which other preferences could be rational. Likewise, if certain states of affairs are of basic final value, then others are not. The basic final value of some states of affairs "constrains" which other states of affairs could be of basic final value.

I thus suggest the following constraint:

Intuitions Constraint: An account of the distinction between basic and non-basic value should confirm our intuitions about what kinds of states of affairs are of basic and non-basic final value.

To be sure, the terms 'basic final value' and 'non-basic final value' are not ordinary terms like 'knowledge' or 'virtue' and so we are less likely to have many robust intuitions about them. But basic final value is meant to be the kind of value that explains why other states of affairs are of final value. We can have intuitions about which states of affairs could have that sort of value, including intuitions about structural or formal constraints.

The *Intuitions Constraint* is formulated in terms of *confirming* our intuitions, not merely being *consistent* with our intuitions. This is intentional. All else being equal, an account of some phenomenon that confirms our intuitions about that phenomenon is preferable to an account that is merely consistent with them. For instance, intuitively, I know that I am sitting right now. An account of knowledge that confirmed that intuition, instead of merely being consistent with it, would be preferable. This point will be important below.

III. Problematic Accounts

In this section, I briefly cover several attempts to draw the basic/non-basic distinction. If we take non-basic final value to be final value that is not basic, then drawing this distinction amounts to giving an account of basic value. I'll briefly argue that several of these accounts are problematic; nonetheless, some of them will point the way forward to a more promising account that I'll sketch in section IV.

A. Quinn's Account

Warren Quinn (1974: 128-32) provides an account for computing value that includes a discussion of basic value. Though Quinn does not have very much to say about basic value, I discuss his account because other authors have, in essence, accepted Quinn's account of basic value and disputed his proposed computations for non-basic value (see, e.g., Oldfield (1977: 234), Carlson (1997: 95-7)).

For Quinn, "basic final value" is the final value had by "basic propositions," which thereby determined "the value of every other proposition" (1974: 128). These basic propositions are not truth-functionally complex (1974: 128). Further, he writes, "It is natural to suppose that the most evaluatively prior of all states of affairs"—i.e. the basic propositions—"are those which locate a specific sentient individual at a specific point along an evaluatively relevant dimension such as happiness, virtue, wisdom, etc." (1974: 131). We might summarize Quinn's proposal like this: the only states of affairs that are of basic final value or disvalue are logically simple, i.e. atomic, states of affairs that attribute to a sentient individual some evaluatively relevant property.

First, on Quinn's account, the states of affairs of basic final value are those that attribute some "evaluatively relevant dimension" to individuals. But Quinn does not describe what exactly makes some property or feature an "evaluatively relevant" one. It cannot simply be that a property is an evaluatively relevant one if it is relevant to whether a state of affairs has final value; for both states of affairs of basic and non-basic final value are evaluatively relevant in that sense. It would be better to say more about this concept and its relation to basic final value. My account, below, will improve on these points.

Second, and relatedly, Quinn's account does not confirm *Intuitions Constraint*. Specifically, Quinn's account does not imply or explain any conditional or comparative claims. Thus, nothing on Quinn's account rules out that each of these are of basic final value:

- (4) At time interval *I*, George experiences some pleasure.
- (5) At time interval *I*, George experiences some pleasure in his beer.
- (6) At time interval *I*, George experiences a great deal of pleasure in his beer.

But, intuitively, these shouldn't all count as being of basic final value.

To be sure, Quinn's account doesn't *imply* that each is of basic final value either. The problem is that it doesn't confirm what is antecedently plausible. Furthermore, there is a straightforward diagnosis of why Quinn's account does not confirm the *Intuitions Constraint*. Quinn's account uses the concept of an "evaluatively relevant" property. But Quinn does not give an account of that concept. And it would not be helpful to say that an evaluatively relevant property is one that is had only by things of basic final value.

I'll briefly mention a final problem: Quinn's proposal is overly strict about the logical form of states of affairs of basic final value. That is, it is conceivable that the states of affairs of basic final value are not always logically simple. For instance, Ross (1930) and Feldman (1995) suggest a view on which the final value of pleasure and pain is partly a function of the degree to which that pleasure or pain is deserved or undeserved. It may be that the best way to regiment such views is in terms of conjunctive states of affairs. Quinn's proposal would unnecessarily rule out such views (or formulation) of such views.

Quinn may be thinking of basic final value as the final value had by a certain type of state of affairs—namely "basic" or non-logically complex states of affairs. That is the wrong way to think about basic final value. Instead, basic final value should be thought of as a type of final value that states of affairs have—namely explanatorily "basic" value that explains the value of other things.

B. Feldman's Account

Feldman's (2000) offers another influential account of basic final value with some authors (e.g., Oliveira (2016), Maguire (2016)) simply adopting his account. However, Feldman does not list necessary and sufficient conditions for something to be of basic final value. He thinks that if he were to list such conditions, it would amount to identifying an actual axiology; but he prefers an account that is neutral—does not beg any questions—about which axiology is correct (2000: 327). Rather, he lists several features such states of affairs would have, regardless of first-order axiological views.

For my discussion, two features are most relevant. The first relevant feature is this: “every basic is a pure attribution of a core [finally] valuable property or relation” (2000: 328). Feldman clarifies what a “pure attribution” of a monadic property is by writing this: “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)” (2000: 328). The second feature is that states of affairs of basic final value are ones that contain no “extraneous” information while giving all the “essential” information (1986: 30; 2000: 328-9).

Ultimately, my own account will owe a lot to Feldman's. But there are still some places for improvement. First, and most importantly, Feldman claims that states of affairs of basic final value contain only “essential” information. But information essential for what? Feldman doesn't clearly specify. Of course, such information is essential for identifying what is of basic final value. But, being true by definition, that answer is trivial and unenlightening. My own account, below, will improve on this point by claiming that the information is relevant to the appropriateness of adopted attitudes.

Second, and relatedly, Feldman does not describe or characterize what a core finally valuable property or relation is. Thus, suppose some property P is a core finally valuable property or relation. Nothing Feldman writes would explain why some other property, P*, would or would not be a core finally valuable property or relation. But sometimes knowing that P is a core finally valuable property *would* inform us that some other property P* is not. For instance, suppose P* is a much less determinate property of P. Then if P is a core finally valuable property, presumably P* wouldn't be.

An alternative way of putting this second complaint is that Feldman's view violates the *Intuition Constraint*. Recall this sequence:

- (4) At time interval *I*, George experience some pleasure.
- (5) At time interval *I*, George experience some pleasure in his beer.
- (6) At time interval *I*, George experiences a great deal of pleasure in his beer.

Suppose that (6) is of basic final value. Intuitively, if (6) is of basic final value, then (4) is not. But nothing Feldman writes explains that intuitive conditional. To be sure, on Feldman's view, if (6) is of basic final value, then it contains only “essential” information. But, for all he's said, perhaps (4) also contains only “essential” information.

What is missing from Feldman's view is the idea that when a property provides “essential” information, then other properties that are related to it may not provide “essential” information. In other words, the “essential” information contained in a state of affairs doesn't float free from the property attributed by the state of affairs. My proposal below will incorporate these points.

C. Zimmerman's Account.

In two works (1999, 2001), Michael Zimmerman offers an account of basic final value. In his (1999: 665) paper, Zimmerman offers this account:

S has basic [final] value =df. S has [final] value, and no proper part of S has [final] value.

This definition uses the concept ‘final value’ and ‘proper part’ to define the concept ‘basic final value.’ In his (2001) book, Zimmerman proposes to inverse this order. After two refinements, he offers a more complex account of basic final value free from the concept of ‘final value’:

Basic [Final] Value 3: S has basic [final] value to a certain degree = df. (a) S is such that the contemplation of it as such directly morally requires that one either favor it to a precisely corresponding degree or disfavor it to a precisely corresponding degree or be indifferent toward it for its own sake and (b) no proper part of S is such that the contemplation of it as such directly morally requires such an attitude, unless that part is a purely temporal proper part of S.

(2001: 161)

Zimmerman then goes on to give an account of final value in terms of basic final value, though the details don’t matter here. What does matter is a corollary of his account:

Zimmerman’s Corollary: S is of basic final value to a certain degree only if (a) S is of final value to that degree and (b) no proper part of S is of any final value.

This corollary follows straightforwardly from his (1999) account. It also follows from his (2001) account, though the details are a little more complicated.⁷

Zimmerman’s account, specifically *Zimmerman’s Corollary*, easily confirms some intuitions about basic final value. For instance, if A is of final value and B is neutral, then it straightforwardly implies that A&B is not of basic final value. For A&B violates requirement (b) of *Zimmerman’s Corollary* that A&B not have any parts of final value. However, as I’ll now argue, it is plausible that sometimes a state of affairs is of basic final value and yet contains a part that is of final value.

My argument turns on the following two claims. The first is:

Less Determinate Value: If a state of affairs S attributing property P is of basic final value, then there is a state of affairs S* that is of final value that differs from S only in that it attributes P* instead of P, where P* is a less determinate property of P.

This claim uses the term ‘less determinate property of’ which I define as:

For any two properties P₁ and P₂, P₁ is “a less determinate property of” P₂ if and only if (i) necessarily, whatever has P₂ has P₁, (ii) possibly, something has P₁ without having P₂, and (iii) P₂ is a way of being P₁.

To illustrate, the property of being red is a less determinate property of being crimson. Similarly, being a basketball player is a less determinate property of being a good basketball player; after all, being a good basketball player is a way of being a basketball player. I will not offer an account of when one property is a “way of being” another, leaving this at the intuitive level.

I lack the space for a full defense of *Less Determinate Value*. Let me just briefly note the following reason for thinking it is true. Whatever states of affairs are of basic final value are likely to attribute very determinate properties. They will presumably be properties that apply to

⁷ In brief, suppose S is of basic final value. From *Basic Final Value 3 (b)* it follows that we are not required to adopt any pro-attitudes towards any proper parts of S. Thus, from *Basic Final Value 3 (a)*, it follows that no proper part of S is of basic final value. Suppose some proper part of S, S*, is of final value but not basic final value. On Zimmerman’s definition of final value (2001: 122), S* contains some parts that are of basic final value. Given the transitivity of parthood, it would then follow that S contains a part that is of basic final value, which we already know cannot happen. Thus, no proper part of S can be of non-basic final value. Since all final value is basic or non-basic, if S is of basic final value, then no proper part of S is of any final value.

particular individuals, are indexed to times, etc. Further, from the fact that such states of affairs are of basic final value it will follow that it is appropriate to value some instances of those very determinate properties. It is plausible that slightly less determinate properties of those properties will also be such that it is appropriate to value some instances of them. For instance, suppose the following state of affairs is of basic final value:

(8) Tully acted kindly from a stable character trait, with an amount of benevolence that was proportionate to the act.

Presumably, if it is appropriate to value (8), it would likewise be appropriate to value:

(9) Tully acted kindly from a stable character trait, with some amount of benevolence.

Or to switch examples suppose the following is of basic final value:

(10) For duration D, Dayanara experienced pleasure in state of affairs S of intensity I, where that amount of pleasure was proportional to the value of S.

If it is appropriate to value (10), it would presumably also be appropriate to value:

(11) For duration D, Dayanara experienced pleasure in state of affairs S, where that amount of pleasure was proportional to the value of S.

Thus, if a valuable state of affairs S attributes a property P to object *o*, then plausibly there is some property Q such that Q is a less determinate property of P and the state of affairs attributing Q to *o* is also of final value. But that is just what *Less Determinate Value* says.

The second claim my argument needs is:

Less Determinate Parts: If two states of affairs S and S* differ only in that S* attributes a less determinate property than S does, then S* is a part of S.

Less Determinate Parts is plausible on the face of it. It is implied by Chisholm's (1986) theory of parthood for states of affairs. On it, S* is a part of S if S's obtaining entails S* obtains and one cannot entertain S without entertaining S*. Plausibly, one cannot entertain a state of affairs S attributing a property P without entertaining the state of affairs S* that differs from S only in that it attributes a less determinate property than P. I cannot entertain the state of affairs of Andrew being a good football player without entertaining the state of affairs of Andrew being a football player. Zimmerman (2001) proposes a different theory of parthood for states of affairs. But it is plausible that on his account *Less Determinate Parts* will hold. On his view, when the instantiation of a less determinate property is explained by the instantiation of a more determinate property, then the particular state of affairs attributing that less determinate property will be a part of the particular state of affairs attributing that more determinate property.

My basic criticism of Zimmerman's view is that it's plausible that something can be of basic final value and yet have a part of final value. Since *Zimmerman's Corollary* implies otherwise, Zimmerman's view is false. More specifically, suppose *Less Determinate Parts* and *Less Determinate Value* are both true. If S is of basic final value, then by *Less Determinate Value* there is a state of affairs that is of final value that differs from S only in that it attributes a less determinate property than S does. From *Less Determinate Parts*, it follows that this state of affairs is also a part of S. Thus, if S is of basic final value, then it has a part that is of final value. Since these two principles are plausible, we should reject *Zimmerman's Corollary* and, by extension, Zimmerman's account of basic final value.⁸

IV. Towards An Account of Basic Final Value

In this section, I'll sketch an alternative account of basic final value, explain the ways in which it handles our earlier problems without falling into the pitfalls of other theorists', before

⁸ For a fuller development of this criticism of Zimmerman, see Perrine (2018).

responding to some objections. The account will build on some of the above accounts, specifically Feldman's and Zimmerman's.

When a state of affairs is of final value, it attributes a property to something. I will call these properties "value conferring properties" (VCPs) because some of their instances confer value. Among value conferring properties I will mark a distinction between "core" value conferring properties (CVCPs) and "non-core" value conferring properties (non-CVCPs). Among the information available in value conferring properties, CVCPs contain the "right" amount of information whereas non-CVCPs do not. Specifically, among the information available in value conferring properties, non-CVCPs can fail to contain the right amount of information by *either* containing irrelevant information *or* failing to include relevant information. A CVCP is thus a value conferring property that neither contains irrelevant information nor fails to include relevant information.

In speaking of "relevant information" I have in mind information that is relevant to the range of pro-attitudes that are appropriate to adopt. One way of specifying which information is relevant is to actually propose an axiology that identifies specific information as relevant or irrelevant. However, since I intend my account of basic final value to be neutral with regard to most axiologies, I cannot pursue such a proposal. Instead, my approach will be to give a general account of when a value conferring property contains irrelevant information or lacks relevant information in terms of the relationships between value conferring properties.

When a non-CVCP contains irrelevant information, that information makes no difference to what range of pro-attitudes are appropriate to adopt. To give an example, compare:

(6) At time interval *I*, George experiences a great deal of pleasure in his beer.

(12) At time interval *I*, George experiences a great deal of pleasure in his beer, where *I* occurred on a Tuesday.

The second state of affairs contains irrelevant information, specifically, that *I* occurred on a Tuesday. The idea that a value conferring property *P* contains irrelevant information can be understood as: there is some other value conferring property *P** such that *P** contains less information than *P* and yet the range of pro-attitudes that are appropriate to adopt towards instances of *P* and *P** are the same. Thus, looking at these two states of affairs, it would be appropriate to adopt the same range of pro-attitudes towards both, even though the second contains more information. In this way, the additional information of the second is irrelevant.

When a non-CVCP lacks relevant information, there is information it does not have that would make a difference to what pro-attitudes are appropriate to adopt. To give an example, compare:

(5) At time interval *I*, George experiences some pleasure in his beer.

(6) At time interval *I*, George experiences a great deal of pleasure in his beer.

Ostensibly, (5) and (6) are about the same event. But (6) includes more information than (5). Further, that difference in information makes a difference to what pro-attitudes are relevant to adopt. Specifically, some general attitude of being pleased may be appropriate upon learning that George experienced *some* pleasure; but a more specific or determinate attitude of being pleased to a greater amount may be appropriate upon learning that George experienced *great* pleasure. In this case, learning more information *refines* the specific range of attitudes that are appropriate. But it is also possible that learning more information would result in adopting entirely contrary attitudes. For instance, consider these states of affairs:

(13) Throughout his life, Ma had many pleasant experiences and few painful ones.

(14) Throughout his life, Ma had many pleasant experiences and few painful ones, where these experiences were caused through his vicious character.

Some philosophers might claim that while (13) might be of final value, (14) is not. They might claim that having a vicious character does not merit pleasure, but pain, and so (14) is of final disvalue.⁹ In such a case, learning more information does not *refine* the specific range of attitudes that are appropriate; rather, it would make appropriate a different set of attitudes that may or may not overlap at all with those that are appropriate for (13). Again, generalizing, the idea that a value conferring property P lacks relevant information can be understood as: there is some other value conferring property P* such that P* contains more information than P and the range of attitudes that are appropriate to adopt towards instances of P* are different from those that are appropriate to adopt towards instances of P. Perhaps the attitudes appropriate to adopt to P* are more specific than those appropriate to adopt towards P; or perhaps they support other attitudes altogether. Either way, P* contains more information than P and that information is more relevant than P's information.

When a property is a value conferring property that is either a more determinate or less determinate property of a core value conferring property, it is natural to see that property as a non-CVCP. A value conferring property that is less determinate than a CVCP will lack relevant information—information that CVCP has. Likewise, a value conferring property that is more determinate than a CVCP will contain irrelevant information—information the CVCP lacks. In this way, it is natural to see a CVCP as being in the middle of a sliding scale composed of various value conferring properties. On the left-hand are non-CVCPs that lack relevant information. On the right-hand are non-CVCPs that contain irrelevant information. As we move from left to right, we move through value conferring properties that contain increasing amounts of information where, somewhere in the center, is a CVCP such that to the left of it are value conferring properties lacking relevant information and to the right of it are value conferring properties that contain irrelevant information.¹⁰

With all this spadework complete, we can now turn to my account of basic final value:

A state of affairs is of basic final value if and only if (a) it is of final value, (b) it attributes, to some particular thing or things, a core value conferring property and (c) there is no proper part of it that attributes, to some particular thing or things, a core value conferring property.

Let us now look at the ways in which this account handles the problems raised above while avoiding some of the pitfalls of other accounts.

First, by drawing a distinction between basic and non-basic final value, this account has the desired ontological flexibility. It permits that some states of affairs—like possible worlds, outcomes of actions, etc.—can be of final value as well as more particular states of affairs.

Second, this account supports the intuition that some cases of over-counting are cases of over-counting. Recall our first case of counting, where $W = A \& B \& C$, A was of final value, but neither B nor C was of any final value or disvalue. To count the value of A&B, A&C, and A when computing the value of W would be, intuitively, to count the value of A too many times.

⁹ Ross (1930: 138) thought that considerations of merit of goods and evil, like pleasure and pain, were relevant to their final value and disvalue. But Ross doesn't explicitly endorse the view I describe here.

¹⁰ Value conferring properties are not always less or more determinate properties of a CVCP. For a value conferring property might simultaneously lack relevant information while also containing irrelevant information. (Consider, for instance, the property of experiencing *some* pleasure in a beer *on a Tuesday*.) In other words, while there are always non-CVCP on either side of the CVCP, not all non-CVCP will fit on a sliding scale with a CVCP.

Given that A is of basic final value, this account implies that neither A&B nor A&C are of basic final value. For both A&B and A&C will contain, as parts, something that attributes a core value property, namely, A. Thus, to count the value of A&B and A&C along with the value of A is a mistake since, given that A is of basic final value, it partly or even fully explains the non-basic value of A&B and A&C.

Our second case of over counting involved these states of affairs:

(4) At time interval *I*, George experiences some pleasure.

(5) At time interval *I*, George experiences some pleasure in his beer.

(6) At time interval *I*, George experiences a great deal of pleasure in his beer.

I claimed that if (6) were of basic final value, then neither (4) nor (5) would be of basic final value. This account secures that result. (4) attributes a less determinate property than (5) and (5) a less determinate property than (6). But any value conferring property that is less determinate than a core value conferring property is not a core value conferring property. Thus, it follows that if (6) attributes a core value conferring property, then (4) and (5) do not. Thus, if (6) is of basic final value, then neither is (4) nor (5).

Third, this account helps avoid over counting involved with computing the instrumental value of things. The resolution of this problem is essentially the same as the resolution of the previous case of over counting only applied to causal consequences.

Finally, this account avoids some of the pitfalls of the other accounts mentioned earlier. First, I objected to both Quinn's account and Feldman's account that they utilized certain concepts in their theory but did not develop accounts of those concepts ("evaluatively relevant" property and "essential information" respectively). Following Feldman, my account uses the concept of "information." But by exploiting the connection between value and appropriateness of attitudes described in section I, I developed an account of relevant information in terms of the appropriateness of adopting attitudes.

Second, and relatedly, I objected to both Quinn's account and Feldman's account that they did not help explain why there would be over-counting when it comes to (4)-(6). Specifically, neither gave resources to explain why if (6) is of basic final value, then neither (4) nor (5) is. My account does have the resources to explain this. If (6) is of basic final value, then it contains a CVCP. Thus, other less determinate properties are non-CVCP. Thus, (4) and (5), which contain a non-CVCP, are not of basic final value. Third, I objected to Zimmerman's account that, intuitively, some things of basic final value have parts that are of final value. My account is consistent with states of affairs of basic final value having parts of final value, so long as those parts are of *non*-basic final value.

Lastly, notice that this account is consistent with "value monism." This is because value monism should *not* be understood as the view that there exists some property P such that whenever a state of affairs is of final value, it attributes P and just P to something. Any view that takes seriously the above points about ontological flexibility will be inconsistent with such a position. Rather, value monism is the position that there is only one core, final value conferring property so that for any state of affairs of basic final value it attributes just P. Value monists need not deny that lots of things are of final value; it is just that such value ultimately derives from one sort of thing. (And, for what it is worth, this is how one value monist—Fred Feldman—understands his value monism; cf. Feldman (2004: 185f).)

V. Information, Valuing, and Partiality

In providing my account of basic final value, I criticized other authors for suggesting that some states of affairs had all the "essential" or "relevant" information without clarifying further

what this information was essential for or relevant to. On my account, such information is relevant to the appropriateness or fittingness of adopting pro- and con-attitudes. However, this account is bound to produce further questions about *when* the information is relevant to adopting attitudes and *which* agents the information is relevant to. In this final section, I briefly explore both of these issues in order to further fill in the elements of my account that are distinctive.¹¹

A. Valuing for Its Own Sake

Earlier, I proposed the following principles:

Final Value: ϕ is of final value if and only if it is appropriate for agents to bear pro-attitudes towards ϕ for its own sake.

Final Disvalue: ϕ is of final disvalue if and only if it is appropriate for an agent to bear con-attitudes toward ϕ for its own sake.

These principles can be seen as refinements of more general principles:

Value: ϕ is of value if and only if it is appropriate for agents to bear pro-attitudes towards ϕ .

Disvalue: ϕ is of disvalue if and only if it is appropriate for an agent to bear con-attitudes toward ϕ .

On this way of thinking, part of what distinguishes things of final value from things of value more generally is the *kind* of attitude that is appropriate for agents to bear. Specifically, in cases of final value and disvalue, certain kinds of attitudes are appropriate or fitting to adopt for their own sake. In this subsection, I'll spend some time describing different approaches to distinguishing between these attitudes before sketching my preferred approach.

One approach is to focus on the *object* of the attitude, which in this discussion are states of affairs. One might propose that only states of affairs about certain things, or with certain information, can be the object of pro- and con-attitudes for their own sake. For instance, perhaps the (mythic) psychological hedonist might propose that people only adopt pro-attitudes towards states of affairs about pleasure for their own sake and only adopt con-attitudes towards states of affairs about pain for their own sake. To illustrate, an instance of this proposal might be:

- For any agent S, S bears pro-attitude towards x for its own sake only if there is a property P such that x is about something having P.

Different concrete proposals might identify different properties. However, this kind of approach is implausible. To be sure, human beings are restricted in what states of affairs they can adopt pro- and con-attitudes towards in virtue of their cognitive limitations. But beyond such restrictions, human beings are able to adopt pro- and con-attitudes towards all sorts of things. Human beings are a pretty varied lot. So it is unlikely that the best way to understand adopting pro- and con-attitudes for their own sake is in terms of some sort of elite property of the objects of those attitudes.

A more plausible approach for understanding bearing pro-attitudes for their own sakes is to characterize such attitudes in terms of bearing pro-attitudes for the sake of other things. There are several ways one might do this. The simplest proposal might be this:

- S bears pro-attitudes towards x for its own sake if and only if (i) S bears pro-attitudes towards x , and (ii) it is not the case that S bears pro-attitudes towards x for the sake of y (where y and x are distinct).

However, this simple proposal has counterintuitive results for unsophisticated agents. My cat Mala bears pro-attitudes towards things. And presumably Mala satisfies clause (ii) simply in virtue of being a cognitive unsophisticated agent who doesn't have pro-attitudes towards things

¹¹ Thanks to an anonymous reviewer for pressing me to address some of these issues.

for the sake of other things. But it is implausible that Mala bears pro- (or con-) attitudes towards things for their own sake.

An alternative proposal might appeal to the structure of the agent's attitudes. When an agent values something for the sake of something else, the agent's attitudes normally do not go on forever. Normally, there is an attitude that is the terminus of the chain. Consequently, one might add a further clause as follows:

- S bears pro-attitudes towards x for its own sake if and only if (i) S bears pro-attitudes towards x , (ii) S bears a pro attitude towards z for the sake of x , and (iii) it is not the case that S bears pro-attitudes towards x for the sake of y (where y and x are distinct).

On this proposal, bearing pro-attitudes towards x for its own sake is equivalent to the pro-attitudes that are at the terminus of bearing pro-attitudes for the sake of other things. In this way, being a pro-attitude towards x for its own sake is the analog of the foundationalist's basic belief.

This proposal will avoid the problem of the last proposal. For this proposal will only apply to agents that have enough cognitive sophistication to bear attitudes towards things for the sake of other things. But this proposal has an obvious problem as well—a problem the previous proposal also suffers from. It is entirely possible for agents to bear pro-attitudes towards something *both* for its own sake *and* for the sake of something else. One might try more complicated proposals of bearing pro-attitudes for its own sake in terms of bearing a pro-attitude for the sake of other things. But these preliminary results cast some doubt that such an approach will succeed.¹²

I propose that these attitudes should not be understood in terms of the content of the attitude or their place in the overall structure of an agent's pro- and con-attitudes. Rather, these attitudes should be understood in terms of the stimulus conditions under which agents are disposed to have the attitudes in the first place. Attitudes have conditions under which agents acquire them; these conditions are frequently called 'stimulus conditions' (for discussion see Miller (2013: chp. 1)). Different kinds of attitudes will have different kinds of stimulus conditions. The stimulus conditions for fear are different from the stimulus conditions for adoration. And there may also be further variation at individual levels. The stimulus conditions under which you and I enjoy jokes might be pretty similar without being identical.

One kind of stimulus condition is contemplating or considering states of affairs in abstraction from other states of affairs. The idea of abstraction is well-known from the sciences. A scientific theory or model might "abstract" from some feature when it doesn't represent that feature or theory in the model; it is "silent" about that feature. (See Cartwright (1989: 187ff.) for a standard discussion of abstraction in the sciences.) So too when one considers a state of affairs in abstraction, one considers the state of affairs abstracted from potential relations to other states of affairs. Though they don't use the term 'abstraction,' I take it that other authors have it in mind. For instance, Chisholm describes a person contemplating something "as such" where contemplating "as such" doesn't include: "the contemplation of some wider situation that one may think that the given object brings along with it" (1986: 52). Similarly, Lemos writes that one contemplates states of affairs "in isolation from the contemplation and ranking of other, wider states of affairs" (1994: 10). Both authors are describing abstraction.

¹² For more sophisticated accounts of this kind, see Fumerton (1990: 132-2), Rønnow-Rasmussen (2022: chp. 8), and Tenen (Forthcoming). For critical discussion, see Perrine (2023).

Of course, one might wonder *what* states of affairs one is abstracting from when one considers a state of affairs in abstraction. A natural answer is the states of affairs one abstracts from are both its causal consequences and the ones it is a part of. In other words:

- An agent A considers or contemplates a state of affairs S in abstraction from other states of affairs only if A does not consider or contemplate other states of affairs, S*, such that either (i) S is a proper part of S* or (ii) S* is a causal consequence of S.

One might propose a further condition:

- An agent A considers or contemplates a state of affairs S in abstraction from other states of affairs only if A does not consider or contemplate other states of affairs, S*, such that S* is a proper part of S.

But this further condition is too stringent. When one considers a conjunctive state of affairs, surely one considers its conjuncts. And even in the sciences, when one has an abstracted model or theory, that model or theory still have parts that are considered in abstraction.

Agents can adopt pro- and con-attitudes in a range of stimulus conditions. One kind of stimulus condition in which they might adopt an attitude is when they consider a state of affairs in abstraction. I propose understanding adopting attitudes for their own sake in terms of what situations are included in the stimulus conditions. Specifically, I propose understanding these attitudes as follows:

- An agent adopts a pro-attitude towards a state of affairs S for its own sake only if the stimulus conditions for their adopting that attitude includes considering S in abstraction.
- An agent adopts a con-attitude towards a state of affairs S for its own sake only if the stimulus conditions for their adopting that attitude includes considering S in abstraction.

Note that these proposals don't assume that there is any uniform feature or property that all states of affairs have when they are the object of attitudes for their own sake. These proposals *do* require agents to be able to consider states of affairs in abstraction, which arguably an unsophisticated agent like my cat Mala lacks. Lastly, these proposals do not exclude agents from adopting pro- or con-attitudes towards something both for its own sake and for the sake of something else. For the stimulus conditions for an agent adopting the attitude might include *both* considering it in abstraction *and* considering it in relation to other states of affairs. In such a case, the agent can rightly be said to adopt attitudes both for its own sake and for the sake of other things.

States of affairs of value attribute properties to things. Those properties contain information relevant to what attitudes are appropriate to adopt towards those states of affairs. Additionally, when states of affairs are of final value or disvalue it is appropriate to adopt those attitudes for their own sake. One important question is *when* this information is relevant to the adopting of attitudes. In this section, I've proposed that the information is relevant when one is considering states of affairs in abstraction from other states of affairs. The next section discusses *who* the information is relevant to.

B. Partiality

As a matter of fact, many human beings show preferences in their pro- and con-attitudes towards people who are close to them in some way, such as family members, coworkers, or fellow compatriots. A common thought is that such preferences are also appropriate or fitting. That is, many people have "partiality intuitions" that considerations of partiality can play a role in determining which attitudes are appropriate or correct for agents to adopt.

In my discussion of final value, I have emphasized the idea that VCP contain information that is relevant to the adoption of pro- and con-attitudes. A natural question is *who* this information is relevant to. Specifically, what role, if any, does partiality play in determining the appropriateness or fittingness of attitudes for agents considering states of affairs with a VCP. While this is a large topic, I can only gesture at some proposals here.¹³

One view is that partial relations are, themselves, value-conferring properties. On this view, information about partial relations is relevant to the adoption of pro- and con-attitudes. Thus, consider these states of affairs:

(15) At time *t*, Danielle helps Filippo.

(16) At time *t*, Danielle, Filippo's father, helps Filippo.

Some might claim that (16) has more final value than (15) exactly because of the information that Danielle and Filippo stand in some partial relation, here father-son relation. These sorts of views are consistent with my account of basic final value, since they simply amount to the claim that some partial relations are VCP. On these proposals, partial relations are just one further property that each of us has reason to care about.

However, some might have intuitions about partiality that support different views than the view that partial relations are VCPs. Rather, they might have the intuition that issues involving partiality ground *variations* in which attitudes are appropriate for which agents. There are at least two kinds of ways partiality might ground variations; and there are at least two kinds of variations that might be grounded. It might be that information about partiality that agents have partly ground variations in attitudes; or it might be that facts about partiality, regardless of whether agents are aware of them, partly ground variations in attitudes. And, regarding variation, it may be that, for some agents, pro-attitudes towards some state of affairs are appropriate whereas for other agents con-attitudes towards that same state of affairs are appropriate. Alternatively, it may be that for some agents more intense pro- or con-attitudes are appropriate than for other agents.

In discussing these kinds of views, I'll take as my foil the following principle:

Uniformity. For any state of affairs with a VCP, the information that state of affairs has fully determines the range of appropriate attitudes for any agent who considers that state of affairs in abstraction.

Uniformity is consistent with the view that partial relations are VCP, that considerations of partiality provide relevant information to adopting pro- and con-attitudes. It is also consistent with the view that it is appropriate for different agents to adopt different attitudes towards a state of affairs in virtue of its consequences. But *Uniformity* is inconsistent with certain views that fill in partiality intuitions. Specifically, it is inconsistent with the view that different agents considering the same state of affairs in abstraction have different ranges of appropriate attitudes.

I take *Uniformity* as my foil partly because I think its plausible. But another motivation for taking it as my foil is that it provides a natural way to block a contradiction. Consider the following state of affairs:

(17) At time *t*, Monique failed to achieve a good, life-long goal of hers.

Now suppose two things. First, that it is appropriate for some agents to adopt con-attitudes towards this state of affairs for their own sake. (Perhaps it is appropriate in virtue of the information contained in the state of affairs or in virtue of their relation to Monique.) Second, it is appropriate for some agent to adopt pro-attitudes towards this state of affairs for their own

¹³ For related discussions, see Blanshard (1961: 287-9), Lemos (1994: 16-18; 2011); Olson (2009); Bykvist (2009); Zimmerman (2001: 119-123; 2011: 472-480), Tappolet (2016:105-110), McHugh and Way (2022).

sake. (Again, perhaps it is appropriate in virtue of the information contained in the state of affairs or in virtue of their relation to Monique.) From the first supposition and *Final Disvalue* (or a principle like it), it follows that this state of affairs is of final disvalue. From the second supposition and *Final Value* (or a principle like it), it follows that this state of affairs is of final value. But that's impossible. Nothing is of both final value and final disvalue. So we have a contradiction. One way of blocking this contradiction, while retaining principles like *Final Value* and *Final Disvalue*, is to accept *Uniformity*. Given *Uniformity*, the two suppositions just made cannot be jointly true.

By contrast, some authors will reject *Uniformity* thereby allowing that it could be possible for the information in (17) to partly make it appropriate for different agents to adopt pro- and con-attitudes towards it. In that case, in order to block this contradiction, those authors will have to reject *Final Value* and *Final Disvalue*. Specifically, those authors will have to reject the claim that the appropriateness of adopting pro- and con-attitudes for their own sake is sufficient for final value and disvalue. Some might claim that the appropriateness of such attitudes is not sufficient for final value or disvalue because there is no such thing as final value or disvalue. As indicated above, I've said aside that view here. Rather, even those who accept the existence of final value and disvalue might claim that the appropriateness of some agents adopting pro- and con-attitudes for their own sakes is insufficient for final value and disvalue. I'll briefly consider two proposals of this kind from Noah Lemos (2011) and Michael Zimmerman (2011).¹⁴

Lemos formulates his views of final value in terms of ethical requirement as opposed to fittingness or appropriateness as I have here. Simplifying, Lemos proposes that when agents consider states of affairs like (17) in abstraction, they have a *prima facie* ethical requirement to adopt con-attitudes towards it for its own sake. However, an agent's *overall* ethical requirement might be to adopt other attitudes other than the ones that are *prima facie* required. Thus, it may be that both a stranger and myself have a *prima facie* ethical requirement to adopt con-attitudes towards this state of affairs. But I have an *overall* ethical requirement to adopt stronger con-attitudes in virtue of partial relations I bear towards my sister that a stranger does not. Thus, the appropriateness of adopting pro- and con-attitudes is not sufficient for final value or disvalue, since what's appropriate for agents may vary whereas the final value or disvalue of the target states of affairs does not.

I'll raise one problem with Lemos' proposal. The problem stems from Lemos usage of the property of ethical requirements as opposed to fittingness or appropriateness. Ethical requirements incur responsibility. If an agent is required to adopt certain attitudes, they are responsible for adopting those attitudes. Further, on voluntarist conceptions of responsibility, if they are responsible for adopting those attitudes, then it is under their voluntary control to do so. But it is doubtful that agents have voluntary control over the attitudes they adopt. So Lemos' proposal may incorrectly restrict the range of value.

Lemos might suggest that, in such cases, agents are required to adopt attitudes towards states of affairs, but if agents are unable to do so, then their inability "defeats" their requirement. However, this understanding of defeat is implausible. If an agent's obligation or duty to ϕ is

¹⁴ Some might posit the existence of so-called agent-relative value. It might be that the appropriateness of adopting pro- and con-attitudes is sufficient for agent-relative value, even if it is not sufficient for final value. But notice that the mere existence of agent-relative value, on its own, wouldn't block this contradiction. One would still have to reject or modify some of these principles—presumably *Final Value* and *Final Disvalue*. For discussions of agent-relative value, see Smith, (2003), Portmore (2005), Schroeder (2007), Wallace (2010: 519ff.), Cullity (2015), Johnson (2021).

defeated, then the agent is still open to evaluations of obligation and duty *vis-à-vis* ϕ -ing; it is merely that the evaluation is no longer one of obligation or duty as opposed to permissibility. However, on a standard view, if an agent is not responsible for ϕ -ing, then the agent is no longer a legitimate target for deontic evaluations. Thus, if the agent is not responsible for ϕ -ing, then the agent doesn't have an obligation or duty that is defeated. Rather, the agent lacks an obligation or duty altogether. So while the distinction between *prima facie* and *ultimate facie* obligations may allow one to distinguish between the scope of agents' obligations, it doesn't straightforwardly apply in this case.¹⁵

Zimmermann (2001: 119-123; 2011: 475-479) offers a different response.¹⁶ Zimmerman accepts:

Basic Uniformity. For any state of affairs with a CVCP, the information that state of affairs has fully determines the range of appropriate attitudes for any agent who considers that state of affairs in abstraction.

But he doesn't accept the more general principle *Uniformity*. That is, he maintains that when it comes to states of affairs of basic final value, then the range of appropriate attitudes is the same for any suitable situated agent. However, it may be the case for certain states of affairs of *non*-basic final value, the range of appropriate attitudes varies from agent to agent. So consider:

(17) At time *t*, Monique failed to achieve a good, life-long goal of hers.

(18) At time *t*, Sarah failed to achieve a good, life-long goal of hers.

If we suppose that these are of basic final value, then the range of appropriate attitudes for me and others is the same. But if we consider:

(19) At time *t*, Monique, who is my sister, failed to achieve a good, life-long goal of hers.

(20) At time *t*, Sarah, who is a stranger to me, failed to achieve a good, life-long goal of hers.

These states of affairs would not be of basic final value. For the fact that Monique is my sister and Sarah a stranger is "evaluatively superfluous" (2011: 475). As a result, the range of attitudes appropriate to bear can vary for me.

I'll briefly raise three reasons why I am dissatisfied with Zimmerman's proposals. First, Zimmerman's proposal is incomplete. There are many properties that are "evaluatively superfluous" to the value of states of affairs. Zimmerman is proposing that amongst some evaluatively superfluous properties some *do* make a difference to what is appropriate to value but not to what is valuable. But Zimmerman doesn't explain or gesture at what, exactly, demarcates them. Second, Zimmerman's proposal is explanatorily defective. Zimmerman proposes, in effect, that for some states of affairs, the range of appropriate attitudes is fully determined or constrained by those states of affairs. But for other states of affairs, the range of appropriate attitudes is not fully determined or constrained by those states of affairs. Dividing explanations in this way is, to me, theoretically unsatisfying. Third, Zimmerman's proposal may be unappealing to those who have strong partiality intuitions. Zimmerman's proposal allows for divergent responses between divergent individuals because the relevant states of affairs contain

¹⁵ Additionally, appropriateness or fittingness don't seem emendable to the *prima facie/ultimate facie* distinction. So one couldn't salvage Lemos' view simply by replacing the property of obligation with these other properties.

¹⁶ McHugh and Way (2022: 254-256) offer a superficially similar response to Zimmerman. However, it is not clear that they are discussing adopting attitudes towards states of affairs. So it's not clear that their response falls within the scope of this work.

irrelevant information about what is valuable. It seems unusual to try to retain the importance of certain intuitions about valuing by claiming they are appropriate responses to information irrelevant to value.

Without offering a full account here, I propose to retain both principles like *Final Value* and *Final Disvalue* with a principle like *Uniformity* while still accounting to some degree for the role of partiality. However, my proposal is fundamentally different from some of the views discussed above. Some of those views try to capture the importance of partiality in terms of intersubjective *variance* of attitudes. By contrast, my proposal tries to capture the importance of partiality in terms of the *volume* of attitudes. In short, for people who bear partial relations to others, there are *more* states of affairs about those others that they consider or contemplate which are appropriate for them to adopt pro- and con-attitudes towards.

Frequently, when people bear partial relations to each other, they know a lot more about each others' lives. That is, when people bear partial relations to each other there's a larger volume of information available to them about each other's lives. The information may include reference to a larger range of events as well as providing more detailed information about those events. As a result, the number of states of affairs about a person's life that one considers or contemplates is normally larger when one bears a partial relation to them. Consequently, the number of pro- and con-attitudes about a person's life are normally much greater when one bears a partial relation to them. Additionally, many of the attitudes are more specific since one knows more specific information about the person's life. Thus, normally when one bears a partial relation to someone, it will be the case that one bears a larger number of more specific attitudes towards states of affairs about that person as opposed to strangers.

This proposal may seem implausible to some philosophers. It may seem implausible due to the frequent idealization that philosophers (myself included!) are prone to use. For instance, consider a common philosophical thought experiment: Two people are drowning and you can only save one. One of the people drowning is a stranger; the other is your daughter. It is stipulated that saving either live will lead to the same, or comparable, outcomes. It is then claimed that intuitively it is more appropriate or fitting to be concerned with the drowning of one's daughter than that of a stranger's.

This case is extremely idealized. In most real life cases, the amount of information one has about one's daughter and her life will far outstrip the amount of information one has about the stranger's daughter. Thus, in most real life cases, considering a calamity like the drowning of one's daughter will bring with it contemplating a huge range of states of affairs that are of final disvalue, basic or otherwise. Because one is contemplating a larger range of states of affairs of disvalue, it is appropriate to adopt a larger volume of con-attitudes.

This proposal in terms of available information and volume of attitudes may also explain other puzzling phenomenon. For instance, Blanshard (1961: 288) points out that oftentimes people think that people bear different and weaker attitudes towards past historical events, even though there is no reason for supposing that suffering of the past was fundamentally different from suffering of the present. And even people who think partiality grounds variations in appropriate attitudes rarely think that *every* partial relation generates appropriate variations. For instance, a child estranged from his father need not bear any special attitudes towards his father. In cases like these, it is plausible that at least part of the reason for why there is variation in attitudes is that there is variation in the amount of information available to the agents. But, again, I don't have the space here to offer a systematic discussion of these puzzling phenomenon.

VI. Conclusion

In this paper, I've developed a theory of basic final value. Central to the theory is the idea that states of affairs have value conferring properties that provide information that is relevant to value. I've developed the idea of relevancy in terms of the range of attitudes that are appropriate for agents to adopt. I've concluded by explaining the situations under which it is appropriate to adopt these attitudes as well as to which agents the information is relevant to.¹⁷

Bibliography

- Adams, Robert. (2007). *A Theory of Virtue*. Oxford University Press.
- Anderson, Elizabeth. (1993). *Value in Ethics and Economics*. Harvard University Press.
- Audi, Robert. (2004). *The Good in the Right*. Princeton University Press.
- Blanshard, Brand (1961). *Reason and Goodness*. Routledge Press.
- Bradley, Ben. (2006). "Two Conceptions of Intrinsic Value." *Ethical Theory and Moral Practice*. 9: 111-30.
- Brentano, Franz (1902). *The Origin of Our Knowledge of Right and Wrong*. Trans. Cecil Hague. Archibald Constable & Co.
- Bykvist, Krister (2009). "No Good Fit." *Mind*. 118.469: 1-30.
- Carlson, Erik. (1997). "The Intrinsic Value of Non-Basic States of Affairs." *Philosophical Studies*. 85.1: 95-107.
- Cartwright, Nancy. (1989). *Nature's Capacities and Their Measurement*. Oxford University Press.
- Chisholm, Roderick. (1976). *Person and Object*. Open Court Publishing.
- Chisholm, Roderick. (1986). *Brentano and Intrinsic Value*. Cambridge University Press.
- Cullity, Garrett. (2015). "Neutral and Relative Value." In *The Oxford Handbook of Value Theory*, eds. Iwao Hirose and Jonas Olson. Oxford University Press.
- Feldman, Fred. (1986). *Doing the Best We Can*. D. Reidel Publishing Company.
- Feldman, Fred. (1995). "Adjusting Utility for Justice: A Consequentialist Reply to the Objection from Justice." *Philosophy and Phenomenological Research*. 55.3: 567-85.
- Feldman, Fred. (2000). "Basic Intrinsic Value." *Philosophical Studies*. 99.3: 319-46.
- Feldman, Fred. (2004). *Pleasure and the Good Life*. Oxford University Press.
- Fumerton, Richard. (1990). *Reason and Morality*. Cornell University Press.
- Gibbard, Allan. (1973). "Doing No More Harm than Good." *Philosophical Studies*. 24.3: 158-73.
- Harman, Gilbert. (1967). "Toward a Theory of Intrinsic Value." *Journal of Philosophy*. 64.23: 792-804.
- Hurka, Thomas (1992). *Virtue, Vice, and Value*. Cambridge University Press.

¹⁷ For helpful feedback and discussion of these issues, I thank David Fisher, Hao Hong, Noah Lemos, Nick Montgomery, Timothy O'Connor, Harrison Waldo, Phil Woodward, Michael Zimmerman as well as several anonymous reviewers for the journal.

- Johnson, Christa (2021). "Leaving Agent-Relative Value Behind." *Utilitas*. 33: 53-67.
- Kagan, Shelly. (1998). "Rethinking Intrinsic Value." *Journal of Ethics*. 2: 277-97.
- Korsgaard, Christine. (1983). "Two Distinctions in Goodness." *Philosophical Review*. 92.2: 169-95.
- Langton, Rea. (2007). "Objective and Unconditioned Value." *Philosophical Review*. 116.2: 157-185.
- Lemos, Noah. (1994). *Intrinsic Value*. Cambridge University Press.
- Lemos, Noah. (2011). "Intrinsic Value and the Partiality Problem." *Philosophy and Phenomenological Research*. 82.3: 697-716.
- Maguire, Barry. (2016). "The Value Based Theory of Reasons." *Ergo*. 3.9: 233-62.
- McHugh, Connor and Jonathan Way. (2022). "Value and Idiosyncratic Fitting Attitudes." In *Fittingness*, eds. Christopher Howard and R. A. Rowland. Oxford University Press.
- Miller, Christian. (2013). *Moral Character: An Empirical Theory*. Oxford University Press.
- Moore, G. E. (1912). *Ethics*. Clarendon Press.
- Nozick, Robert. (1981). *Philosophical Explanations*. Harvard University Press.
- Oldfield, Edward. (1977). "An Approach to a Theory of Intrinsic Value." *Philosophical Studies*. 32.3: 233-249.
- Oliveira, Luis R. G. (2016). "Rossian Totalism about Intrinsic Value." *Philosophical Studies*. 173: 2069-86.
- Olson, Jonas. (2009). "Fitting Attitude Analyses of Value and the Partiality Challenge." *Ethical Theory and Moral Practice*. 12: 365-378.
- Perrine, Timothy. (2018). "Basic Final Value and Zimmerman's *The Nature of Intrinsic Value*," *Ethical Theory and Moral Practice*. 21.4: 979-996.
- Perrine, Timothy. (2023). "Is it Fitting to Divide Value? A Review of the *Value Gap*." *Journal of Moral Philosophy*. 20: 533-544.
- Portmore, Douglass. (2005). "Combining Teleological Ethics with Evaluator Relativism." *Pacific Philosophical Quarterly*. 86: 95-113.
- Rabinowicz, Wlodek and Toni Rønnow-Rasmussen. (2000). "A Distinction in Value: Intrinsic and for Its Own Sake." *Proceedings of the Aristotelian Society*. 100: 33-51.
- Rabinowicz, Wlodek and Toni Rønnow-Rasmussen. (2003). "Tropic of Value." *Philosophy and Phenomenological Research*. 66.2: 389-403.
- Rønnow-Rasmussen, Toni. (2022). *The Value Gap*. Oxford University Press.
- Ross, W. D. (1930). *The Right and the Good*. Clarendon Press.
- Ross, W. D. (1939). *The Foundations of Ethics*. Clarendon Press.
- Quinn, Warren. (1974). "Theories of Intrinsic Value." *American Philosophical Quarterly*. 11.2: 123-32.

- Scanlon, T. H. (1998). *What We Owe to Each Other*. Harvard University Press.
- Schroeder, Mark. (2007). "Teleology, Agent-Relative Value, and 'Good.'" *Ethics*. 117.2: 265-295.
- Smith, Michael. (2003). "Neutral and Relative Value After Moore." *Ethics*. 113.3: 576-598.
- Tappolet, Christine (2016). *Emotions, Values, and Agency*. Oxford University Press.
- Tenen, Levi. (Forthcoming). "How Final and Non-Final Valuing Differ." *Journal of Ethics*.
- Tucker, Miles. (2016). "The Pen, The Dress, and the Coat: A Confusion in Goodness." *Philosophical Studies*. 173.7: 1911-1922.
- Wallace, R. J. (2010). "Reasons, Value, and Agent-Relativity." *Dialectica*. 64.4: 503-528.
- Zimmerman, Michael. (1999). "Virtual Intrinsic Value and the Principle of Organic Unities." *Philosophy and Phenomenological Research*. 59.3: 653-66.
- Zimmerman, Michael. (2001). *The Nature of Intrinsic Value*. Rowman and Littlefield Publishing
- Zimmerman, Michael. (2010). "Responsibility, Reaction, and Value." *Journal of Ethics*. 14: 103-15.
- Zimmerman, Michael. (2011). "Partiality and Intrinsic Value." *Mind*. 120.478: 447-483.