

Thinking Globally, Acting Locally: Partiality, Preferences and Perspective

Graham Oddie

Volume 9, Number 2, Summer 2014

URI: <https://id.erudit.org/iderudit/1026678ar>

DOI: <https://doi.org/10.7202/1026678ar>

[See table of contents](#)

Publisher(s)

Centre de recherche en éthique de l'Université de Montréal

ISSN

1718-9977 (digital)

[Explore this journal](#)

Cite this article

Oddie, G. (2014). Thinking Globally, Acting Locally: Partiality, Preferences and Perspective. *Les ateliers de l'éthique / The Ethics Forum*, 9(2), 57–81.
<https://doi.org/10.7202/1026678ar>

Article abstract

A rather promising value theory for environmental philosophers combines the well-known fitting attitude (FA) account of value with the rather less well-known account of value as *richness*. If the value of an entity is proportional to its degree of richness (which has been cashed out in terms of *unified complexity* and *organic unity*), then since natural entities, such as species or ecosystems, exhibit varying degrees of richness quite independently of what we happen to feel about them, they also possess differing degrees of mind-independent and subject-independent value. In particular, their value is not dependent on the desires or preferences of humans. The fitting attitudes account of value, at least as it is standardly developed, demands isomorphic evaluative responses on the part of all valuers. In particular, it entails that all valuers should have isomorphic preferences. But this seems absurd. I consider three different strategies with which the fitting attitude theorist can deflect this challenge. The first makes use of an account of non-standard value relations in terms of permissible preference orderings. The second appeals to value appearances and the associated notions of value distance and value perspective. The third involves an account of the ultimate bearers of value as properties, rather than as propositions or states of affairs. These strategies are not all mutually incompatible. While it isn't possible to combine the first and second strategies, it is possible to combine the first and third strategies, and also to combine the second and the third.



THINKING GLOBALLY, ACTING LOCALLY: PARTIALITY, PREFERENCES AND PERSPECTIVE

GRAHAM ODDIE
UNIVERSITY OF COLORADO

ABSTRACT:

A rather promising value theory for environmental philosophers combines the well-known fitting attitude (FA) account of value with the rather less well-known account of value as *richness*. If the value of an entity is proportional to its degree of richness (which has been cashed out in terms of *unified complexity* and *organic unity*), then since natural entities, such as species or ecosystems, exhibit varying degrees of richness quite independently of what we happen to feel about them, they also possess differing degrees of mind-independent and subject-independent value. In particular, their value is not dependent on the desires or preferences of humans. The fitting attitudes account of value, at least as it is standardly developed, demands isomorphic evaluative responses on the part of all valuers. In particular, it entails that all valuers should have isomorphic preferences. But this seems absurd. I consider three different strategies with which the fitting attitude theorist can deflect this challenge. The first makes use of an account of non-standard value relations in terms of permissible preference orderings. The second appeals to value appearances and the associated notions of value distance and value perspective. The third involves an account of the ultimate bearers of value as properties, rather than as propositions or states of affairs. These strategies are not all mutually incompatible. While it isn't possible to combine the first and second strategies, it is possible to combine the first and third strategies, and also to combine the second and the third.

RÉSUMÉ :

Pour les philosophes de l'environnement, une théorie des valeurs plutôt prometteuse allie l'analyse de la valeur en termes d'attitudes appropriées, déjà bien connue, à une analyse moins bien connue, celle de la valeur en termes de *richesse*. Si la valeur d'une entité est proportionnelle à son degré de richesse (celle-ci étant rendue en termes de *complexité unifiée* et d'*unité organique*), étant donné que les entités naturelles, comme les espèces ou les écosystèmes, présentent des degrés de richesse variables aucunement liés à ce que nous pourrions ressentir à leur égard, elles possèdent également divers degrés de valeur indépendante du sujet et de la pensée. En particulier, leur valeur ne dépend ni des désirs ni des préférences des êtres humains. L'analyse de la valeur en termes d'attitudes appropriées, tout au moins selon son élaboration habituelle, exige des réponses évaluatives isomorphes de la part de toute entité attribuant une valeur. Or, cela suppose que ces dernières doivent posséder des préférences isomorphes, ce qui nous semble absurde. Nous examinons trois différentes stratégies que peut adopter le théoricien de l'attitude appropriée pour faire face à ce défi. La première met à profit une analyse non standard des relations de valeur en termes de classements de préférence permis. La deuxième fait appel aux apparences de valeur et aux concepts connexes de distance de valeur et de perspective de valeur. La troisième fait recours à l'analyse des porteurs ultimes de valeur en tant que propriétés, et non comme propositions ou états de fait. Ces trois stratégies ne sont pas toutes mutuellement incompatibles. Bien qu'il soit impossible de combiner la première et la deuxième stratégie, il est possible d'unir la troisième à la première et à la deuxième.

FITTING ATTITUDES AND VALUE AS RICHNESS

A fitting attitude (FA) account of value deems something valuable just to the extent that it is appropriate, or fitting, or obligatory to take a certain attitude towards it. The account of value as richness holds that the value of an entity is proportional to its degree of *richness* (or its *unified complexity* or its *organic unity*). The combination of these two theses yields a promising value theory for environment philosophers. Natural entities, such as species or ecosystems, exhibit varying degrees of richness, and they do so quite independently of what we happen to feel about them. So, if value is richness, these entities also possess differing amounts of mind-independent and subject-independent value. Their value is not dependent on the desires or preferences of humans. And given the fitting attitudes account, their value makes certain affective and conative responses on the part of humans and other receptive beings fitting, appropriate, or obligatory. So this combination entails that whatever our actual responses to the natural world happen to be, it is fitting to embrace, preserve, pursue and enjoy the value of those rich systems that populate it¹.

However, the fitting attitude account of value faces a partiality problem. The account, at least as it is standardly developed, demands isomorphic evaluative responses from all valuers. If something is good, say, then every valuer should favour it. And if one thing is better than another, then every valuer should prefer the former to the latter. The value-responses of all valuers should thus be isomorphic to the value facts and hence to each other. Consequently, if two valuers have differing preferences, then at least one of them must be responding unfittingly. But that, it seems to me, is quite incredible. There appear, to me at least, to be large areas over which responses to value not only do, but should vary from one valuer to another.

I will consider three different ways for the fitting attitude theorist to avoid the isomorphic response challenge. The first makes use of Rabinowicz's account of non-standard value relations in terms of permissible preference orderings. The second appeals to value appearances and the associated notions of value distance and value perspective. The third involves an account of the ultimate bearers of value as properties, rather than as propositions or states of affairs. These different strategies are not all incompatible, and in the end, the fitting attitude theorist may have to invoke more than one of them in order to accommodate apparently legitimate variations in our preferences and responses to value.

FITTING ATTITUDES AND THE ISOMORPHIC RESPONSE OBJECTION

The fitting attitude schema ties values to attitudes in a systematically pleasing way. At first blush, the schema seems rather obvious in the case of certain thick evaluative attributes: the admirable is what it is fitting or appropriate to admire, the revolting what it is fitting or appropriate to be revolted by, and so on. But it is also at least plausible for the thin evaluative attributes: goodness seems closely connected to what is desirable, and the desirable is what it is fitting to desire².

If the FA schema is applicable to all values, then it holds out the promise of a reduction of values to their associated attitudes together with the notion of fittingness. Where *favouring* and *disfavouring* are placeholders for the fitting response to bearers of V , the fitting attitude schema is this:

X has value/disvalue V if and only if it is fitting to favour/disfavour X .

A fitting attitude reductionist would strengthen the schema: entity X has value/disvalue V if and only if *and because* it is fitting to favour/disfavour X . This strengthening is intended to capture the idea that the value is possessed by the entity *in virtue of* the fittingness of the attitude, not vice versa. But this reductionist bid founders on the well-known and much discussed Wrong Kinds of Reasons objection³.

Suppose we begin with a deontic notion of fittingness—i.e., favouring is fitting if favouring is the attitude that you *ought* to have or adopt. Now suppose an evil demon threatens the world with something terrible unless you admire him for making that very threat. Then you ought to admire him, even though he isn't in fact admirable⁴. I will come back to what I take to be the fitting attitude theorist's best response to the WRK problem; what I want to examine first is a much less discussed objection to the FA schema.

The basic FA schema matches up values with attitudes. As standardly articulated, the fittingness of a response is not supposed to be a subject-relative matter. It is fitting for one to favour X if and only if it is fitting for everyone to do so. It follows that the responses of any two valuers who are responding fittingly (call these "fitting responders") are isomorphic with the value facts and hence isomorphic with each other. Now suppose, as seems plausible, that desire is the fitting response to the desirable (the good) and preference is the fitting response to betterness. If one valuer prefers X to Y and another prefers Y to X , then one or the other (or possibly both) must be responding unfittingly. But that seems a tad counterintuitive. Surely there are some occasions that invite a range of different preferences, all of which are fitting responses to value. Different valuers can, quite legitimately, respond in different ways to things which possess their values independently of those valuers' responses.

Here is an example. Consider Erin. Erin loves egrets. She has always loved egrets. Egrets are her passion. She grew up near some wetlands (hereafter "the Wetlands"), spending all the spare time in her childhood mucking about in them—studying, photographing, drawing and monitoring the egrets and the other wildlife there—and she still does this whenever she can. She studied evolutionary biology at college, but after that went into a career in equity, and now has a lucrative day job working for a hedge fund. The Wetlands have recently been threatened by a housing and supermarket development project that will completely destroy the egrets' habitat along with the rest of the rich ecosystem in which they thrive. Distressingly for Erin, it is the very hedge fund for which she works that has a controlling share in the project. Erin is deeply conflicted—she

loves her job otherwise and makes a huge amount of money—but eventually she decides to leave the hedge fund behind, to take much less well remunerated part-time work, and to spend her own time founding the Wetlands Preservation Project to fight the development.

Tammy, some ten thousand miles away, loves the Takahe, a native bird of New Zealand. Only four were observed in the nineteenth century, and by 1930 they were believed extinct. But although they came perilously close, they were not in fact extinct. After a sighting in 1948 in Fiordland, some modest preservation strategies were implemented, and now there are roughly two hundred birds. Tammy is keen to ensure the continuation of the Takahe in their habitat, but she has no particular aptitude for field work or environmental activism and decides that the best way for her to promote the survival of the Takahe is to work harder at her lucrative career in real estate and devote all the extra cash she earns to the Takahe Preservation Foundation.

Now suppose Erin also learns about the Takahe project. And suppose she has an opportunity to re-join the hedge fund, where her apostasy would be forgiven. She could go back, abandon the Wetlands to her firm's development project, and send all the extra money she makes to the Takahe Preservation Foundation. But she just has no real desire to do that, so she sticks with her modest part-time work and continues pouring her energies into the Wetlands Preservation Project.

Tammy and Erin appear to have different desires and preferences. Erin prefers doing hands-on work with the Wetlands project to working harder in her lucrative career and donating cash to the Takahe project. Tammy, by contrast, prefers to work harder at her lucrative career in real estate and donate her spare cash to the Takahe Foundation. Do these differences entail that at least one of these two preference structures is unfitting? Is Erin's pursuit of the Wetlands preservation project together with her decision not to return to her lucrative job and pursue the Takahe preservation project less fitting than Tammy's opposite and equally enthusiastic response? Is Erin's continuing commitment to the Wetlands project—even after she learns of the Takahe project—misguided? Or is Tammy's response defective? Should she get out of real estate and, like Erin, pitch in at the ground level?

Erin and Tammy's different ways of responding to these values both seem quite legitimate or fitting. Any value theory stating that at least one of these two valuers must be responding unfittingly—on the grounds that, to respond fittingly, both must exhibit isomorphic responses to the same value bearers—seems wrong. In particular, a value theory that implies that Erin can only respond fittingly by changing her preferences in some way—by abandoning her desire to work on the Wetlands project, to work harder at her job on Wall Street, and devote her excess monetary resources to the Takahe project—seems highly questionable, if not absurd.

VALUE AS RICHNESS

On the face of it, there are very different kinds of valuable things: happiness, desire satisfaction, pleasure, knowledge, virtues, works of art, music, persons, justice, living organisms, species and ecosystems, to name a few. Recently, Chris Kelly (2003 and 2014) has argued that a certain theory of value can capture most of these very diverse value judgments⁵. It is the thesis that value just is, or is proportional to, the *richness*, or *unified complexity*, or *organic unity* of the value bearer.

Value as Richness (VR)

Degree of value is proportional to degree of richness (or unified complexity, or organic unity).

This is a monistic theory of value. It has had a small number of subscribers in the history of philosophy (most notably Leibniz and Nozick⁶), but it is neither well known nor fully developed yet. It promises to capture, unify and explain a number of otherwise disparate intuitions about value. In particular, it makes rather good sense of the judgments that many environmental philosophers have urged concerning the mind- and subject-independent value of entities such as ecosystems. Indeed, ecosystems are paradigm exemplars of rich entities, for they exhibit high degrees of unified complexity or organic unity.

As a monistic theory value, VR is compatible with a strong form of value comparability.

Strong Comparability

Any two entities (of the same ontological type) are strongly comparable with respect to their degree of unified complexity (or degree of richness): either *X* is richer than *Y*, or *Y* is richer than *X*, or *X* and *Y* are equally rich.

Now suppose, as also seems plausible, that the fitting response to betterness is *preference*. Then we have:

Preferences should track richness.

One ought to prefer *X* to *Y* if and only if *X* is richer than *Y*.

Consider then the relative richness of Erin's Wetlands ecosystem and Tammy's Takahe ecosystem. Assuming strong comparability, there is some fact of the matter about which of the two is richer. If the Takahe ecosystem is richer, then both Erin and Tammy ought to prefer the preservation of the Takahe ecosystem to the Wetlands ecosystem (all other things being equal). If the two ecosystems have an exactly equal degree of richness, then Erin and Tammy should both be indifferent between them. Presumably they should also be indifferent as to which one to commit resources to, deciding perhaps by the toss of a coin. If the two ecosystems are of exactly equal richness, then any tiny increase in the richness

of one ecosystem should move both conservationists to line up behind that preservation project. So, for example, a tiny increase in the value of pursuing the Takahe ecosystem should be enough to move Erin to abandon her commitment to the Wetlands ecosystem. But this seems odd. Any account of the fittingness of responses should be able to accommodate diverse legitimate responses, including those exhibited by Tammy and Erin in the scenarios above.

Suppose that the preservation of the Takahe ecosystem would make the world richer than would the preservation of Erin's Wetlands. This seems plausible: there are plenty of good egret habitat available worldwide. Even if the development of Erin's Wetlands would eliminate the egret ecosystem from Erin's particular locality, egrets and their supporting ecosystems are under no particular threat globally. The failure of Tammy's Takahe preservation project, on the other hand, would seriously diminish the richness of the world. So, if preferences ought to track richness, then Erin ought to prefer the success of the Takahe preservation project to the success of her Wetlands preservation project. Presumably, a consequence of this theory is that she ought to want to increase her hours in the odious Wall Street hedge fund (the one that is destroying her beloved Wetlands), and send all her spare cash to the Takahe preservation project. But she doesn't have those desires and preferences. She stubbornly prefers to promote the poorer outcome (the preservation of the Wetlands ecosystem) to the richer outcome (the preservation of the Takahe ecosystem). So she is an unfitting responder. One could, of course, simply bite the bullet here and embrace the conclusion that all responders should have isomorphic preferences. But that seems rather extreme, so here I consider three possible ways for the proponent of FA to resist⁷.

PARITY AND PERMISSIBLE PREFERENCES

In his "Value relations" (and in a number of subsequent papers), Wlodek Rabinowicz develops a normative (or deontic) notion of fittingness within a general fitting attitude account of value relations:

X is more valuable than *Y* if and only if it is *deontically* fitting to prefer *X* to *Y*⁸.

He notes that there are two different normative, or deontic, notions of fittingness—a strong notion and a weak notion. A preference might be *obligatory* (the strong notion of fittingness) or it might be *permissible* (the weak notion of fittingness).

Weak deontic fittingness

It is weakly fitting to prefer *X* to *Y* if and only if it is *permissible* to prefer *X* to *Y*.

Strong deontic fittingness

It is strongly fitting to prefer *X* to *Y* if and only if one *ought* to prefer *X* to *Y*.

Rabinowicz formally models the permissibility and obligatoriness of particular preferences in terms of a collection of rationally permissible preference orderings.

It is a common assumption in value theory that both rational preferences and betterness induce complete orderings in their domains. This seems rather a tall order. It would be more reasonable to countenance some gaps in the orderings—that is, to countenance *partial* orderings of the elements in the domain to be compared for value. Allowing partial orderings won't by itself eliminate the isomorphic response objection. For if there is only one permissible partial ordering then only one preference ordering is permissible (albeit an ordering with gaps). And, by the FA schema, that preference ordering would be isomorphic to the betterness relation. In order to avoid the isomorphic response objection, the FA-theorist must therefore countenance more than one permissible value ordering (with or without value gaps in it). Allowing more than one permissible preference ordering yields a certain latitude in what preferences are permissible—the larger the class, the greater the latitude.

Indeed, the model was developed at least in part to accommodate Chang's non-standard value relation of *parity* within a fitting attitude approach. X and Y are on a par if they are in some sense comparable for value, but none of the standard value relations hold (namely: X is better than Y ; Y is better than X ; X has exactly the same value as Y). Chang appeals to the *small improvement argument* to motivate the possibility of parity. For example, consider two ecosystems that seem to be in the same ballpark of value: Erin's Wetlands and, say, a comparably rich grassland ecosystem in Colorado. Is one more valuable than the other? One might be reluctant to subscribe to either judgment. There are of course epistemic problems involved here. There are many elements of both ecosystems to evaluate and weigh, and one might overlook some of them or make a wrong judgment. There is also the problem of how to combine the value of the individual components into an overall evaluation of the ecosystem. If the ecosystem is valuable in proportion to its richness (or organic unity), then its value clearly will not be the sum of the values of individual components. But the problem might be deeper than these epistemic factors. It could be that neither ecosystem is in fact superior to the other, yet it does not follow that they must then be either of equal value or else incomparable for value.

Suppose one grants that neither ecosystem is in fact more valuable than the other. Since richness involves at least two dimensions (unity and complexity), given VR it is possible that neither is clearly better than the other⁹. The Grasslands might exhibit greater unity, while the Wetlands exhibit greater complexity, say, and it may not be determinate how to weight these two factors in an overall assessment of richness. Now consider a small improvement to Erin's Wetlands. Suppose a small quantity of some mildly harmful chemical has been slowly leaching into the Wetlands from a nearby industrial plant, and Erin manages to get this source of contamination sealed off, thereby enhancing the unity of the Wetlands. Call this improved wetland ecosystem Wetlands⁺. Wetlands⁺ is clearly more valuable than Wetlands. And if Wetlands and Grasslands are equal in value, then Wetlands⁺ would also be clearly better than Grasslands. But while Wetlands⁺ is clear-

ly better than Wetlands, that does not seem to make Wetlands⁺ clearly better than Grasslands. So Wetlands and Grasslands aren't equal in value either. But neither are Wetlands and Grasslands totally incomparable for value. They are clearly in the same ballpark. For example, if an environmental agency had a limited amount of money to spend on a substantial improvement to one ecosystem, both a Wetlands proposal and a Grasslands proposal might legitimately make it into the final round, ahead of various other clearly inferior proposals. If Wetlands won and was *substantially* improved thereby (call that Wetlands⁺⁺), then the result might well be richer, and hence more valuable, than Grasslands.

Betterness is given by strong fittingness of preference (one *ought* to prefer X to Y) which holds just in case X is preferred to Y in every permissible preference ordering. X and Y are *on a par* provided that it is (weakly) fitting to prefer X to Y but that it is also (weakly) fitting to prefer Y to X . That is, in some permissible preference orderings, X is preferred to Y , while in others, Y is preferred to X . A small improvement to X (X^+) will be preferred to X in every preference ranking (i.e., X^+ is strictly better than X), but if the improvement is small enough, then X^+ and Y , like X and Y , will also be on a par. However, a substantial enough improvement in X (X^{++}) would be better than Y —i.e., X^{++} is preferred to Y in every permissible preference ranking.

Suppose, then, that we list a number of permissible (kinds of) preference ranking involving our various ecosystems:

Preference ranking A: Takahe $>_A$ Wetlands⁺⁺ $>_A$ Wetlands⁺ $>_A$ Wetlands $>_A$ Grasslands

Preference ranking B: Takahe $>_B$ Wetlands⁺⁺ $>_B$ Wetlands⁺ $>_B$ Grasslands $>_B$ Wetlands

Preference ranking C: Takahe $>_C$ Wetlands⁺⁺ $>_C$ Wetlands⁺ $>_C$ Wetlands \approx_C Grasslands

Preference ranking D: Takahe $>_D$ Wetlands⁺⁺ $>_D$ Wetlands⁺ \approx_D Grasslands $>_D$ Wetlands

.....

In every permissible preference ranking: Takahe is preferred to Wetlands and Takahe is preferred to Grasslands. Also Wetlands⁺⁺ is preferred to Wetlands⁺, which in turn is preferred to Wetlands. Where $>_V$ is the relation of *greater value*, the intersection of all the permissible preference rankings, we have:

Takahe $>_V$ Wetlands⁺⁺
 Wetlands⁺⁺ $>_V$ Wetlands⁺ $>_V$ Wetlands
 Takahe $>_V$ Grasslands.

But none of the following relations hold:

Wetlands $>_V$ Grasslands
 Grasslands $>_V$ Wetlands;
 Wetlands \approx_V Grasslands.

Wetlands and Grasslands are on a par given this collection of admissible preference rankings. Rabinowicz's model thus allows for the fittingness of a preference for Wetlands over Grasslands, or for Grasslands over Wetlands, as well as indifference between the two. All these are fitting preferences—that is, they are all *weakly* fitting. So not all fitting preferences have to be isomorphic either to actual value or to each other.

Despite the latitude that Rabinowicz's model allows, it cannot spare Erin condemnation for her preferring the Wetlands project to the Takahe project. The model cannot countenance the compatibility of Takahe $>_V$ Wetlands with the permissibility of a preference for Wetlands preservation over Takahe preservation. But it is worse than this. Suppose that the Grasslands preservation is neither better nor worse than Wetlands preservation, but exactly the same value. This means that in every permissible preference ranking they are equi-preferred. So there is no permissible preference ranking in which Wetlands preservation is preferred to Grasslands preservation. If this is the case, then Erin's pursuing the Wetlands' preservation over the Grasslands' preservation is *strictly impermissible*. Although she may toss a coin to decide whether to support one cause or the other, if she does launch herself whole-heartedly and passionately into the Wetlands preservation without so much as a second thought for the equally valuable Grasslands and the value of their preservation, then her desires are to that extent defective.

REPRESENTATION, PERSPECTIVE AND DISTANCE

If fittingness is a deontic notion, a matter of what one *ought* to prefer, then the *Wrong Kinds of Reason* objection looms¹⁰. Suppose a powerful but evil demon threatens to bring about the very worst outcome unless you admire him, and will spare the world the worst outcome only if you do so admire him. Then you ought to admire him—even though he is not in the least admirable. Similarly, if the demon threatens to bring about the very worst outcome unless you prefer that outcome to all others, and will spare the worst outcome if you do so prefer it, then you ought to prefer it—but the outcome you ought to prefer is not good. The same kind of counterexample can be made to work against an axiological notion of fittingness as well (i.e., that the good just is what it is good to favour).

Various attempts, mostly rather ad hoc and unsatisfactory, have been made to square these normative readings of fittingness with WKR counterexamples. I will not review those here. But an FA theorist can embrace an alternative notion of fittingness that sidesteps the WKR counterexamples altogether. Tappolet calls it a *descriptive* reading, as opposed to a *normative* reading. She motivates this notion by appealing to a perceptual theory of emotions, whereby an emotion involves a non-conceptual representation of value:

An alternative conception [...] is that the appropriateness of emotions is a matter of representing things as they are. In the relevant sense, appropriate emotions are emotions that are correct from an epistemic point of view¹¹.

I will call this the *representational* (rather than descriptive) notion of fittingness. Where A^V is the attitude that pairs with value V , the FA schema becomes:

X is V if and only if it is representationally fitting for one to take attitude A^V to X .

Suppose that attitude A^V consists in the object X 's *appearing* to have value V . For an attitude A^V to X to be representationally fitting is for A^V to represent X as possessing the value V , and for X to in fact possess the associated value V . For example, for you to admire the aforementioned demon, the demon must appear in some way admirable to you. And it is fitting to admire the demon just in case that appearance represents the demon accurately. Representational FA will not, of course, underwrite a *reduction* of value to non-evaluative concepts, since on this account, the fitting attitudes themselves must have evaluative representational content (i.e., they are representations of things as possessing value). But if one is not intent on reducing goodness to some other properties, then this is not an objection to the representational reading of fittingness as such. A fitting attitude theorist can endorse the schema without touting it as a reductive tool.

The Representational FA theorist can easily avoid the demon objection and others like it. She can say: "It isn't fitting to admire the demon, for that attitude would be a grossly inaccurate representation. It would be to represent the demon as admirable, when he is in fact contemptible. But the best thing to do, all things considered, is to avert the demon's threat of harm, even though that involves taking an attitude towards the demon that clearly doesn't represent him as he is. In order to do the right (or best) thing, you simply have to suppress your (perfectly fitting) contempt for the demon and do your best to admire him." Similarly, even if desire is the right or best attitude to take to the worst outcome in light of the demon's threat, that does not make desire a *fitting* response to the worst outcome *per se*.

How does the representational FA schema bear on the problem of isomorphic preferences? Augustine articulated an early version of what is sometimes called the *guise of the good*:

[...] [In] the [pull] of the will and of love, appears the worth of everything to be sought or to be avoided, to be esteemed of greater or lesser value¹².

I call this the *value appearance* thesis. It has a number of related components. First, there are appearances of value: these are perception-like rather than belief-like (something can *appear* a certain way without one's believing it to *be* that way.) Second, a desire for X is or involves an appearance of the goodness of X (and a preference for X over Y is or involves an appearance of X 's possessing more value than Y .) Third, an appearance of value is fitting just in case it is an accurate representation of its object.

The isomorphic response thesis would accordingly take the following form:

X is better than *Y* if and only if it is fitting for *X* to appear better than *Y* (if and only if it is fitting to prefer *X* to *Y*). So the preferences of any two fitting responders will be isomorphic.

The analogue of this idea for regular perception would be patently absurd. It is not a requirement of accurate perception that for any two observers the appearances should be isomorphic to reality and hence isomorphic to each other. This is because perceptual representations are representations *from a certain point of view*, and as such they legitimately incorporate perspectival effects. Suppose two objects *X* and *Y* are exactly the same size. Dom is closer to *X* than to *Y* while Eric is closer to *Y* than to *X*. To Dom, *X* will appear larger than *Y*, while to Eric, *Y* will appear larger than *X*. But neither response is an unfitting or inaccurate perceptual representation of the objects. Rather, it is fitting that the objects should appear differently to each perceiver given the different relations that the former bear to the latter. Perception is always perception of objects as they stand in relation to the perceiver. From the Earth, the setting moon should appear larger than the rising sun, even though it is in reality much smaller. As such, the appearances for two differently situated perceivers should not in general be isomorphic to reality or to each other.

If there are genuine appearances of value, then there should also be value analogues of distance and perspective. If we suppose that desires and preferences value appearances, then they should depend not just on the evaluative properties of the objects of those appearances but also on how one stands in relation to those objects. There should be a perspectival element to desire, just as there is a perspectival element to visual perception. Just as it is entirely fitting for distant objects to appear small by comparison with nearby objects, states that are far from you in value space should not loom as large in your desires and preferences as states that are close to you. The weals and woes of those who are closely connected to you, those for whom you care deeply, are closer to you in value space than the weals and woes of distant strangers. Suppose, for example, that pain is bad, and that qualitatively identical pains have the same disvalue whoever is the subject of that pain. Suppose that your daughter and some stranger are experiencing extreme and qualitatively identical pains. You are far more averse to your daughter's being in extreme pain than to the stranger's being in pain. Of course you are also averse to the stranger's predicament, but if there were only one dose of morphine at hand, then you would not find yourself completely indifferent between the stranger's getting that dose and your daughter's getting it. But if we take a deontic approach to fittingness, then it is clear that you *ought* to be completely indifferent as to whether the last dose of morphine goes to your daughter or to the stranger. That seems odd. I think that it is completely appropriate or fitting for you to care more about the pain of people you love and with whom you have deep connections than about the pain of people with whom you have no connections at all—which is not to say that it would be fitting not to

care at all about the stranger's pain, nor that you should believe that your daughter's pain is worse. I mean simply that it is appropriate for you to prefer that your daughter receive the last dose of morphine. If desires and preferences are non-doxastic appearances of value then this falls out rather naturally.

We can now apply this perspectival analysis to the conception of value as richness. Some rich states of affairs are much closer in value space to a particular value responder than other equally rich states. The survival of the Wetlands is much closer in value space to Erin than the survival of the Grasslands, even though the two ecosystems are (let's suppose) equally rich and thus equally valuable states of affairs¹³. Even though they are equally valuable, it is entirely fitting that the Wetlands should loom larger in Erin's preferences and desires (that is, in how valuable they appear to her) than the Grasslands. If it is entirely fitting for Erin to prefer the preservation of the Wetlands to the preservation of the Grasslands, then it is also fitting for her to be more motivated to promote the preservation of the Wetlands over that of the Grasslands. And if she acts on such fitting motivations, then presumably that is something for which we cannot fittingly criticize or condemn her.

One might object that this is all well and good in states of affairs that are of *exactly* equal value, or even perhaps for states that are on a par. In Erin's case, her closer connection to the Wetlands can legitimately tip the balance only because there is either a value tie or because there is parity. But if one state of affairs is more valuable than another, then such connections become completely irrelevant. The only fitting response is to prefer the more valuable state to the less valuable one. But if distance can make a difference to legitimate preferences in the case of exact equality of value, then it would constitute an inexplicable discontinuity if it could not make any difference at all in a case of two states that are extremely close in value but not quite equal.

PROPERTIES AS VALUE BEARERS

There are two problems with the perspectival reply to the isomorphic response objection. First, the notion of value perspective and distance in value space, as it stands, is somewhat metaphorical. Is there any way of making it less metaphorical, of theoretically cashing out the notion of perspective? Second, it is not clear why desire and preference should be fitting responses to the (distance-discounted) richness of states of affairs. Why should the richness of states of affairs present itself to a value responder as something that it is fitting to desire?

The value appearance thesis is that desires and preferences are appearances of value. (It is *not* the thesis that all appearances of value are desires.) For the value appearance thesis to be on the cards, the objects of desire must also be bearers of value. The standard view is that desire and preference, like belief, are *propositional* attitudes, and that the objects of such attitudes are either propositions or closely related entities that propositions pick out, namely states of affairs. This dovetails nicely with a widespread monistic view about value bearers, namely that propositions or states of affairs are the fundamental bearers of intrinsic

value¹⁴. If other things—such as individuals or properties—have value, then they have it in virtue of the intrinsic value of the states of affairs in which they feature.

The propositional attitude view sits happily with the surface grammar of some desire claims, but not of all. Consider:

Oliver desires the Olympic gold medal for the 15km cross-country skiing event.

Harvey wants a hokey-pokey ice-cream.

Erin prefers the Wetlands to the Grasslands.

An Olympic gold medal, a hokey pokey ice cream, Wetlands and Grasslands all appear to function here as objects of desire, or of preference, but they are not obviously propositions or states of affairs. However, whenever a desire is specified without specifying an explicit propositional object, it is usually easy to find an essentially equivalent and natural specification of the desire that makes a propositional analysis plausible. It is not too farfetched to say that Oliver's desire for *the Olympic gold medal*, or Harvey's desire for *a hokey-pokey ice-cream* are desires for certain interactions with the entity at issue: to *win the gold* or to *eat a hokey-pokey ice cream*. What Erin prefers is *preserving the Wetlands* to *preserving the Grasslands*, or *working on the Wetlands preservation project* to *working on the Grasslands preservation project*. Finally, Oliver's desire *to win the Olympic gold* will be further parsed as Oliver's wanting it to be the case *that Oliver wins an Olympic gold*. But what state does Erin prefer to what? Erin prefers *that Erin preserve the Wetlands* to *that Erin preserves the Grasslands*. This last formulation may sound a bit strange, but one can get used to it fairly quickly.

There is, however, a rival to the propositional view that secures the uniformity of the ontological category without any cost to naturalness or flexibility. Suppose Oliver and Orlando are rivals for the Olympic gold in the 15km cross-country skiing event. Oliver says *I really want to win Olympic gold!* Orlando chimes in *So do I!* The common object of their two desires is *to win the Olympic gold medal in the 15km cross-country event* (or *to win the Olympic gold* for short). On the propositional view, they can only want the same thing if they want the very same state of affairs to obtain: but there is no state of affairs that corresponds to their common desire in this case. Rather, what each wants is a *property*, that of *winning the Olympic gold*. In general, what one desires is to possess or to acquire some property, and what one prefers is the having of some property to the having of some other property¹⁵. The property view can make literal sense of the claim that Oliver and Orlando want *the very same thing* (to have the property of *winning the Olympic gold*). If Oliver's desire is satisfied, then Orlando's is frustrated, and vice versa.

In addition to wanting to preserve the Wetlands, Erin also wants *the Wetlands to be preserved*. But *the Wetland's being preserved* is a state of the world, rather than

a property of Erin. So, at the very least, the property view has to be able to capture the class of desires having states of the world as their apparent objects. Happily, there are quite natural property-surrogates for states of the world that can serve as the objects of these apparently state-directed desires. I will call these *global* properties. For each state of the world S there is exactly one global property, $P(S)$, characterized as follows: X has property $P(S)$ if and only if S obtains. $P(S)$ is the property of *its being such that S obtains*¹⁶. Since X doesn't occur on the right-hand side of the definition, $P(S)$ is *global* in the sense that it is had by some individual if and only if it is had by all. Either every individual has $P(S)$ or none have it. The circumstances in which it is had by one and all are just those in which the world is in state S . So wanting to have $P(S)$ is tantamount to wanting S itself to obtain. We can retrieve the class of propositional desires, or state desires, from desires the objects of which are these global properties. When Erin and Tammy both have the apparently state-directed desire *that the Wetlands be preserved*, the common object of their desire is a global property.

Global properties can be contrasted with *local* properties. A local property is a property that can be had by some individuals without being had by all. The property of *preserving the Wetlands* is local, since Erin can have it while Tammy lacks it.

What does desire fulfillment consist in, on the property view? On the state or propositional view, the desire that P is fulfilled if and only if the object of the desire, the proposition P itself, is true. But properties do not bear truth values. Further, on the property view, Oliver and Orlando can share the *very same* desire (*winning the Olympic gold*). But how can those be the *same* desire if the fulfillment of Oliver's desire will constitute the frustration of Orlando's? That is a violation of Leibniz's principle. However, this puzzle is easily dissolved. *Desire* suffers from the usual state/object ambiguity: *Oliver's desire* can pick out the state that consists in Oliver's desiring something, or it can pick out the object of Oliver's desiring. Fulfillment is not solely a feature of the object of desire. For Oliver's desire (to win the Olympic gold) to be fulfilled, *Oliver* (rather than Orlando or anyone else) must acquire the object of his desire. In this case, Oliver must acquire the property of winning the gold.

Consider a selection of local and global properties: winning the Olympic gold; eating a hokey pokey ice-cream; the *Goldberg Variations* being played; hearing the *Goldberg Variations* being played; playing the *Goldberg Variations*; preserving the Wetlands; the Wetlands' being preserved by someone or other. These are all possible objects of desire on the property view, but they are not only objects of desire—they are also apt subjects of value attributions. Plausibly, all these properties have some value, although some are more valuable than others. Eating a hokey pokey ice cream has some minor value; hearing the *Goldberg Variations* being played is more valuable; and playing the *Goldberg Variations* is more valuable still. Winning the Olympic gold may be somewhere in the ballpark of playing the *Goldberg Variations*. So on the property view, it is plausible that the possible objects of desire are also apt bearers of value.

Of course, not all bearers of value are possible objects of desire. The *Goldberg Variations* is a valuable musical work. Even if it turns out to be a property (say,

the property had by all and only possible performances of the work), that is not the kind of property that it makes sense to *desire*. No rational person wants to *instantiate* that property. To so desire would be a rather odd category error, as one can rationally desire only those types of property that one could conceivably instantiate. However, this is no obstacle to the value appearance thesis, since the latter does not require that all bearers of value be possible objects of desire.

What makes this shift in the primary bearers of value—from propositions to properties—a rather promising step in tackling the isomorphic preference problem is this: local properties (unlike propositions or global properties) have their own built-in *perspective* or *point of view*. A local property of individuals does not characterize a state of the world; rather it characterizes a way of *being in the world*.

For example, consider the following four properties, each specifiable in terms of the more basic properties of playing the *Goldberg Variations* and listening to the *Goldberg Variations*: *playing the Goldberg Variations to an audience of at least one* (P); *being a member of an audience listening to the Goldberg Variations being played* (H); *neither playing nor listening to the Goldberg Variations* (N); *its being such that the Goldberg Variations are played by one person and heard by at least one other* (G). P, H and N are local properties, G is global. P, H and N characterize three very different kinds of niche in any world in which G is instantiated. If you are told that you (and hence everyone) will come to have the property G, then you thereby acquire some interesting information about the world and its richness. But you would be none the wiser as to whereabouts in the structure of the world you yourself would end up—you might be a *Goldberg Variations* player, or a *Goldberg Variations* listener, or you might be neither. Each of the local properties P, H and N does not simply characterize the G-world, but rather characterizes locations within the structure of a G-world.

More generally, if we take any individual *X*, we can ask: *What basic monadic properties does X have, and what basic monadic properties does X lack?* There is a collection of complete specifications of basic monadic properties which partitions the set of individuals into mutually exclusive and jointly exhaustive kinds. Call this a *depth-0 kind*. We can go on to ask: *What basic relations does X bear to various depth-0 kinds of individuals?* This will refine the initial partition into a more fine-grained partition. These are depth-1 kinds of individuals. Each individual falls into one and only one of these depth-1 kinds. We can go on increasing the characterization of an individual's properties and relations to other individuals to any depth of specificity we like, and as we go deeper and deeper, we narrow down more precisely the particular kind of niche that a given individual occupies in the world. At an infinite depth, we arrive at near complete specifications of the kinds of loci that individuals can occupy in the web of being¹⁷. These infinitely deep kinds are rather close to Leibniz's complete individual concepts, though without further assumptions one cannot show that the niches so specified can only be occupied by one individual, or that an infinitely deep kind can only be instantiated in one possible world.

THE RICHNESS OF PROPERTIES

According to VR, states are valuable in proportion to their richness. States are rich in virtue of their *structure*. To establish this, it is sufficient to note that two isomorphic states of affairs must enjoy the same degree of richness or unified complexity. The best explanation for this relationship is that it is the *structure* of a state that determines its degree of richness. In other words, a state's degree of richness corresponds to the richness of its structure. But the structure of a state is the common *property* of all the states that are isomorphic to it. Thus, states possess richness in virtue of the richness of the properties that they instantiate. Consequently, if the thesis of value as richness is applicable to states, it is also applicable to properties.

Some properties are clearly richer (exhibit a higher degree of unified complexity) than others. *Being a human being* involves having a richer structure than does *being an amoeba*. And *being an amoeba* involves having a richer structure than does *being a disorderly heap of stones*. Moreover, these judgments of relative richness accord with the correlative judgments of relative value. Going back to our examples in the previous section, eating a hokey pokey ice cream has a fairly low (albeit non-negative) degree of richness, while hearing and appreciating the *Goldberg Variations* has a greater degree of richness, and playing the *Goldberg Variations* with appropriate skill and appreciation has greater richness still.

Let E be any ecosystem, and consider two local properties and one global property associated with E . (Here D indicates *direct participation*; I indicates *indirect participation* and N indicates *non participation*.)

E^{Dir} : To preserve ecosystem E , to which one is already richly connected through past interactions, by resigning from one's lucrative job and directly devoting one's time and energy to preservation activities.

E^{Ind} : To preserve ecosystem E by working harder at one's lucrative job and donating one's spare cash to E 's preservation.

E^{Non} : To be such that the E 's ecosystem gets to be preserved.

(Note that property E^{Dir} builds in the individual's existing rich connections to the ecosystem E . The other two properties do not, but neither do they exclude that possibility.)

Erin wants to give up her job at the hedge fund and start the Wetlands Preservation Project. Abbreviating Wetlands to \mathbf{W} , Erin *prefers* \mathbf{W}^{Dir} to \mathbf{W}^{Ind} and she *prefers* \mathbf{W}^{Ind} to \mathbf{W}^{Non} . To Erin, it thus seems more valuable to be directly involved in preserving the Wetlands than to be only indirectly involved (i.e., by staying on

at the hedge fund and donating cash to the cause), which in turn seems more valuable than the Wetlands's being preserved by someone or other. Assuming that value is richness, these appearances could be entirely fitting. \mathbf{W}^{Dir} is (plausibly) a richer property to have than either \mathbf{W}^{Ind} or \mathbf{W}^{Non} . If the Wetlands are preserved by someone or other, then *everyone*, no matter what relations they individually bear to the Wetlands, has the global property \mathbf{W}^{Non} . For Erin to have \mathbf{W}^{Non} , the Wetlands must be preserved, but Erin does not have to bear any other connections to them. But to have \mathbf{W}^{Dir} , Erin must not only have the global property \mathbf{W}^{Non} , but in addition she herself must already have various rich connections with the Wetlands and also bear interesting new causal and intentional relations to their preservation. In particular, if she possesses \mathbf{W}^{Dir} , she will be integral to the complexity of Wetlands in a way that she may well not be if someone else preserves the Wetlands courtesy of infusions of her cash.

We can motivate this idea by comparing the relative richness of another pair of properties: playing a rich piece of music, like the *Goldberg Variations*, and paying someone else to play it for you. Suppose you love the *Goldberg Variations*. To hear the *Goldberg Variations* being played with appreciation is a property with a rather high degree of richness. The *Goldberg Variations* themselves have a high degree of unified complexity, of course, and whoever hears and appreciates them gets to participate in that richness. Paying to have someone else play them so that you can hear them is quite a rich property to have. But one who both appreciates *and* skilfully plays the *Goldberg Variations* not only gets to hear them played: other things being equal, she is much more closely tied to that complex unified structure than is the person who merely pays to have them played, by virtue of her being the direct causal conduit through which that structure becomes realized in the world¹⁸. The mental states of the listener and the player both add richness to the original richness of the work itself, but the player's mental state is more closely united with the work than is the listener's state. So the appreciative and skilful player of the *Goldberg Variations* possesses a richer property than the appreciative but passive listener.

If this analysis is right, it suggests that Erin, by participating directly in the preservation of the ecosystem to which she is already richly connected, would become richer than if she simply donated the cash to someone else to do the preserving of the Wetlands (or of some equally valuable ecosystem to which she bears no connection) in her stead. And that in turn is richer than the Wetlands' getting preserved somehow or other. So:

$$\mathbf{W}^{\text{Dir}} >_{\mathbf{V}} \mathbf{W}^{\text{Ind}} >_{\mathbf{V}} \mathbf{W}^{\text{Non}}.$$

Now suppose the Grasslands (\mathbf{G}) have exactly the *same* degree of richness as the Wetlands. Then even though $\mathbf{W}^{\text{Non}} \approx_{\mathbf{V}} \mathbf{G}^{\text{Non}}$, we have that $\mathbf{W}^{\text{Dir}} >_{\mathbf{V}} \mathbf{W}^{\text{Ind}} >_{\mathbf{V}} \mathbf{G}^{\text{Non}}$. So engaging directly or indirectly in the preservation of an ecosystem will be a richer property than its being such that an equally valuable ecosystem is preserved somehow or other. Now consider a slightly richer Grasslands, \mathbf{G}^+ . If

\mathbf{W}^{Ind} is better than \mathbf{G}^{Non} , then \mathbf{W}^{Dir} might well be better than $(\mathbf{G}^+)^{\text{Non}}$, since the added richness involved in one's directly preserving the Wetlands could exceed the added richness involved in Grasslands⁺. This means that there is no obstacle in principle to the possibility that \mathbf{W}^{Dir} is more valuable even than \mathbf{T}^{Non} (where Takahe = \mathbf{T}). It all depends on the differences in richness involved. Of course, the Takahe ecosystem might be so much richer than the Wetlands that the additional richness in the property \mathbf{W}^{Dir} is swamped by the extra richness of \mathbf{T}^{Non} . In that case, Erin's closeness to and connection with the Wetlands cannot make it fitting for her to prefer directly preserving those to indirectly preserving a much richer ecosystem with which she has no pre-existing deep connection.

Finally, even if participating indirectly in the preservation of the Takahe ecosystem (by staying at her hedge fund and donating the money to the Takahe Preservation Fund) is more valuable than its being such that the Takahe ecosystem is preserved by someone or other (i.e., $\mathbf{T}^{\text{Ind}} >_{\mathbf{V}} \mathbf{T}^{\text{Non}}$), the extra richness of the former may not fully make up for the difference in richness between \mathbf{W}^{Dir} and \mathbf{T}^{Non} . Hence, directly participating in the preservation of the Wetlands may still be a richer property to have than indirectly participating in the Takahe preservation.

So the following forms a consistent set of value relations:

$$\mathbf{W}^{\text{Dir}} >_{\mathbf{V}} \mathbf{T}^{\text{Ind}} >_{\mathbf{V}} \mathbf{T}^{\text{Non}} >_{\mathbf{V}} \mathbf{W}^{\text{Non}}.$$

Assuming these are the actual value relations among these properties—it's my story!—it would be entirely fitting for Erin's preferences to be isomorphic to these, and we can endorse that judgement without having to invoke the notions of value distance and value perspective that we need to invoke when we take states to be value bearers.

Note that Tammy shares some of these preferences. She prefers \mathbf{T}^{Ind} to \mathbf{T}^{Non} and prefers \mathbf{T}^{Non} to \mathbf{W}^{Non} . But Tammy is not at all desirous of giving up *her* lucrative job and participating directly in the preservation of the Wetlands. This does not generate a discrepancy between her and Erin, however. Since Tammy does not bear any pre-existing rich relations to the Wetlands, she is not eligible to possess \mathbf{W}^{Dir} . So her decision to devote her money to the Takahe preservation fund does not show that she prefers \mathbf{T}^{Ind} to \mathbf{W}^{Dir} . Even if she gave up her job in real estate, she could not acquire the property \mathbf{W}^{Dir} by going to work with Erin, since she would simply not have the prerequisites for it.

But what about the fact that Tammy prefers the *indirect* preservation of the Takahe (i.e., by keeping her job and donating the cash) to direct preservation (i.e., by giving up her job and working for the Takahe Foundation)? Is this not a case where her preferences, by failing to align with the real values, are defective? It is possible, but it may be also be that these are not the salient properties for Tammy to be weighing. It may be that she simply cannot possess \mathbf{T}^{Dir} . Recall that Tammy has no aptitude for that kind of work, and so any attempt on her part

to acquire T^{Dir} will end in complete failure. Or it might also be that although Tammy really does prefer T^{Dir} to T^{Ind} , her choice is really T^{Dir} and some property P that is logically stronger and richer than T^{Ind} , a property which implies T^{Ind} , excludes T^{Dir} , and which is much richer than either of those. It might be that by staying on and making loads of cash in real estate she can instantiate many more rich properties (she might, for example, have the opportunity to learn to play the *Goldberg Variations*).

The same sorts of considerations might apply to Erin's preference for W^{Dir} over T^{Dir} . She bears close causal, historical, and epistemic connections to the Wetlands ecosystem by virtue of her many years of interacting with them. By contrast, she bears none of these connections to the Takahe ecosystem, so she cannot acquire T^{Dir} even if she wanted to. Further, even if the Takahe's being preserved, T^{Non} , is a richer global property than the Wetlands' being preserved, W^{Non} , directly preserving the Wetlands, W^{Dir} , is a richer local property than the Takahe's being preserved (T^{Non}) or even indirectly preserving the Takahe (T^{Ind}).

Finally, all of this is compatible with the assumption that the global property of the Takahe's being preserved appears more valuable to Erin than does the Wetlands' being preserved. So Erin can fittingly prefer the local property of directly preserving the Wetlands to the local property of indirectly preserving the Takahe, even though she prefers the state consisting in the Takahe ecosystem's being preserved to the state consisting in the Wetlands ecosystem's being preserved.

The property account of value bearers gives the fitting attitude theorist a wide scope for accounting for apparently differing preferences among valuers without deeming some of them defective and without having to appeal to value distance or value perspective. The work done by value distance and value perspective on the state theory is accomplished instead by the differing degrees of richness of the connections that a valuer can bear to those independently rich systems.

This leaves one last problem. Why should desires line up with the appearances of goodness/richness in the way that the value appearance thesis says they should? This is a particularly stark problem if states of affairs are the basic bearers of value. Suppose you are contemplating two states of affairs and it becomes clear to you that one of them has a slightly higher degree of unified complexity than the other. Why should the former necessarily strike you as *better*? Why should you *ipso facto* prefer the former to the latter?

Suppose an asteroid is heading toward a pair of twin inhabited planets, and it is not clear which one it is going to hit and destroy. The planet you are on (call it *Planet*) has slightly smaller degree of unified complexity than its twin does (call that one *Twin Planet*). When you are huddling with your family waiting for calamity to strike either Planet or Twin Planet, should you prefer that Planet, along with your nearest and dearest, be obliterated lest the universe suffer a slightly greater diminution of overall richness?

Now suppose that properties, rather than states, bear value. If the asteroid hits your planet, most of the properties you will end up with will be very impoverished indeed. You, and all the folk with whom you enjoy rich connections, will be reduced to rubble. That is an extremely poor set of properties to end up instantiating. But if it is fitting for you to prefer richer properties to poorer properties, then it is fitting for you to prefer surviving largely intact to being obliterated, even if this results in some very small overall cost to the overall richness of the world. It is better to be rich than to be poor, and so it is better to survive with a high degree of richness than to be reduced to a heap of rubble¹⁹.

The combination of these three theses—of properties as value bearers, of value as richness, and of preference as the fitting response to betterness—yields the following important consequence: it is entirely fitting for you to prefer the richer local property of *surviving intact along with your nearest and dearest while the other planet is destroyed* (S) to the poorer local property of *being reduced to rubble along with your nearest and dearest while the other planet survives* (R). And your twin on Twin Planet can, entirely fittingly, also prefer S to R. Thus, your preferences can be perfectly isomorphic both to the value relations themselves as well as to each other.

But here is a possible objection. The global property of *Twin Planet's surviving while Planet is reduced to rubble* (TS) is (by assumption) a slightly richer property than the global property of *Planet's surviving while Twin Planet is reduced to rubble* (PS). So it would also be fitting for you to prefer global property TS to global property PS. But then we seem to have an irresolvable conflict between fitting preferences over global properties and fitting preferences over local properties. (There is no such conflict for your more fortunately placed twin.) Possessing the richer of the two local properties, in these circumstances, precludes possessing the richer of the two global properties. Your preferences over local properties and over global properties, despite both being fitting according to the theory, thus come into irresolvable conflict. If your preferences over global properties are eventually satisfied (that is, Planet is reduced to rubble while Twin Planet survives) your preferences over the local properties are not (you, along with your nearest and dearest, become a bloody scattered mess in the rubble, an extremely poor state to end up in) and vice versa. Since not all your fitting preferences can be satisfied, this combination of theses is apparently incoherent.

While much more needs to be said about this objection, I don't think it delivers a fatal blow to the combination. Fitting preferences should track the values. But for almost any theory of value, there will be occasions when not all values can be realized, and so not all fitting preferences can be satisfied either. In this case there are just two ways of combining local value (local richness) with global value (global richness). One possibility combines the realization of vastly greater local richness, but at a tiny cost to global richness; and the other possibility combines an enormous diminution to local richness with a tiny gain in global richness. With that way of framing it, it does not seem totally irrational to prefer the former to the latter, all things considered.

CONCLUDING REMARKS

The Isomorphic Response objection suggests that at least one of our two environmental heroines—Erin or Tammy—is responding unfittingly because their preferences appear not to be isomorphic to each other. Our three strategies all give the fitting attitude theorist who embraces value as richness some resources for resisting the objection and deeming both Tammy and Erin’s responses fitting despite their apparent differences.

The first strategy countenances a range of permissible preferences at the outset. While the betterness relation serves to constrain the fittingness of responses, it does not prescribe one unique preference ordering to all rational valuers. There can thus be multiple different yet legitimate responses to richness, and two fitting responders can have non-isomorphic preferences. But even though this gives responders some latitude in their preferences, it does not permit any responder to prefer one state of affairs to another state of affairs that is equally rich. So it is unfitting for Erin to prefer the preservation of her beloved Wetlands to the preservation of some distant but equally rich ecosystem with which she has no connection. Since these are equally rich ecosystems, she should simply be indifferent between them. And if Erin *thinks* that they are equally rich, then it would be incoherent of her not to be indifferent.

The second strategy takes the notion of value appearances seriously and attributes a perspectival aspect to such appearances. Perceptual representation depends not only on the properties of the perceived objects themselves but on the relation of the perceiver to those objects. Objects that are closer to a perceiver should loom larger for her in her visual perception than objects that are more distant. Analogously, some loci in value space are closer to a responder than are others; accordingly, it is sometimes fitting for one of two equally rich states to *appear* better than the other to a particular responder. So it is fitting for the preservation of the Wetlands to *seem* better to Erin than the preservation of the distant Grasslands. If preferences are value appearances, then it is fitting for Erin to *prefer* the preservation of the Wetlands to the preservation of the Grasslands: even if she *thinks* that the Grasslands are no less rich than her Wetlands, she can nevertheless legitimately ignore the Grasslands and instead work assiduously to preserve her beloved Wetlands.

The third strategy takes properties, rather than states of affairs, to be value bearers. Global properties (those that are had by all if had by one) can do double-duty as states. But there are many local properties that are both objects of desire as well as bearers of value. The value of a property is determined by its richness, just as (on the value as richness account) the value of a state is determined by its richness. For Erin, preserving her beloved Wetlands is a richer property to have than the property of preserving the remote and unloved Grasslands. This is quite compatible with the Wetlands’ being no richer than the Grasslands. Some relations which Erin can bear to the Wetlands are rich not solely in virtue of the complexity of the Wetlands themselves, but also in virtue of her connections to the

Wetlands. And since she bears these relations to the Wetlands but not to the Grasslands, she will become richer by preserving the former rather than the latter. Her preference for one local property over the other is not only fitting (the former is richer than the latter) but is also compatible with her thinking the two ecosystems to be equally rich and equally valuable in themselves.

The fitting attitude theorist is not necessarily forced to choose just one of these strategies. The first and the third seem to be compatible. One might hold that there is more than one permissible preference ordering on the grounds that richness does not induce a single complete ordering of properties. Further, one could combine this position with the idea that some properties are determinately richer by virtue of relations between the property-bearer and other entities that the property involves. One might also combine the second and the third strategies. It may be that directly preserving the Wetlands is preferable to indirectly preserving the Takahe ecosystem due to the richness of the respective properties. But it might also be that even if the former is not a richer property, it should loom larger in Erin's desires simply because the Wetlands are closer to her in value space. It does not seem possible, however, to combine the first and second strategies. For the first requires that Erin be indifferent between equally good properties, whereas the second may require her to prefer one to the other.

What is clear, I think, is that the combination of value as richness, a fitting attitudes account of the representational kind, and properties as value bearers yields a rich array of resources for environmental value theory that deserve further exploration²⁰.

NOTES

- ¹ See Kelly 2014 (this volume) and Kelly 2004 for an exposition of the idea of value as richness and its own rich history. The idea can be found explicitly in both Leibniz and Nozick, but it is also implicit in a lot of work on value theory. One can take the proposal, as Kelly does, to constitute a complete and monistic theory of value, in which goodness is in some sense identified with richness. Or, perhaps more plausibly, one can take it to be a partial account of value, in which richness is one of a number of thick determinates of the thin concept of goodness. For my part, I am inclined to think that the latter is the more plausible, but for the purposes of this paper I can leave this question open.
- ² Rabinowicz 2013.
- ³ Not all fitting attitude enthusiasts are value reductionists by any means. See, for example, Zimmerman 2001 and Wedgwood 2009. The latter argues explicitly that it is a mistake to conflate the universal applicability of the FA biconditional (if it is indeed universal) with the idea that the normative concept of fittingness has some kind conceptual priority with respect to the axiological concepts (p. 516). Further, as I argue below, one can embrace the universal applicability of FA biconditionals while adopting a *non-normative* notion of fittingness—fittingness as representational correctness—and this combination clearly eschews the possibility of reduction altogether. So one can embrace the FA biconditional schema without embracing the idea that the FA schema yields a reductive analysis of value.
- ⁴ Rabinowicz and Rønnow-Rasmussen 2004.
- ⁵ Kelly 2004 and 2014.
- ⁶ See Kelly 2014 : “God has chosen the [rule] which is the most perfect, that is to say the one which is at the same time the simplest in hypotheses and the richest in phenomena, as a geometrical line might be, of which the construction was easy and the properties and effects very admirable and of great extent.” Leibniz 1953, p. 9. “Holding fixed the unifiedness of the material, the degree of organic unity varies directly with the degree of diversity of that material being unified. Holding fixed the degree of diversity of the material, the degree of organic unity varies directly with the degree of unifiedness [...]” Nozick 1981, p. 416.
- ⁷ The Isomorphic Response objection was raised in Oddie 2005. It was also raised by A.C. Ewing as an objection to his own fitting attitudes account. See Olson 2009, pp. 365–378.
- ⁸ Rabinowicz 2008.
- ⁹ See the quote from Nozick 1981, footnote vi.
- ¹⁰ Rabinowicz and Rønnow-Rasmussen 2004.
- ¹¹ Tappolet 2011, p. 119.
- ¹² St. Augustine 1982, p. 109.
- ¹³ For a longer development of the notions of distance in value space and of its perspectival effects, see Oddie 2005, chapter 8.
- ¹⁴ See Zimmerman 2010, section 4.
- ¹⁵ Lewis 1979. Note that one does not have to buy into the baroque aspects of Lewisian metaphysics to find the property view of desire plausible. Indeed, the view is probably more plausible if worlds are abstract ways that things can be, rather than large, causally isolated hunks of junk. Nor does one have to buy into Lewis’s companion thesis that the objects of belief are also properties, though that view sits happily enough with what I am endorsing here.
- ¹⁶ This is not how Lewis characterizes the property view in his 1979 paper, but his characterization is dependent on his idiosyncratic view of individuals and worlds—in particular his denial of any trans-world identity.
- ¹⁷ This talk of depth of specification is not just a metaphor. The depth of a characterization is a measure of the number of nested quantifiers that one needs to use to obtain that characterization. What I have been describing informally here are the depth-*d* distributive normal forms discovered by Hintikka (see his 1963 paper). These have been put to good philosophical work in both inductive logic and truthlikeness. See Oddie 1986 chapter 4 for an accessible introduction.

- ¹⁸ As noted by an anonymous referee, other things may not always be equal. If Bach listened to some amateur playing the *Goldberg Variations*, the relation of Bach's mental state to the played music would doubtless add far greater richness to the world than that added by the relation of the amateur's mental state to the played music.
- ¹⁹ Incidentally, this example also gives us a glimpse of how VR might be combined with the thesis that properties are value bearers to defuse the evolutionary debunking argument against value realism. But that is a topic for another paper.
- ²⁰ I would very like to thank the sponsors and organizers of the one-day workshop on *Attitudes, Values, Environment* at the University of Montreal in February 2014 (particularly Christine Tappolet), and the two anonymous referees for this journal, for many excellent comments and criticisms on earlier versions of this paper.

REFERENCES

St. Augustine, "The Literal Meaning of Genesis," in J. Quasten, W.J. Burghardt and T.C. Lawler (eds.), *Ancient Christian Writers: Vol. 1 Book 4*, New York/Mahwah, NJ, Paulist Press, 1982.

Hintikka, Jaakko, "Distributive normal forms in first-order logic," in J.N. Crossley and M.A.E. Dummett (eds.), *Formal Systems and Recursive Functions (Proceedings of the Eighth Logic Colloquium)*, Amsterdam, North-Holland, 1963, pp. 47-90.

Kelly, Crhis, *A Theory of the Good*, Ph.D. Thesis, University of Colorado, Boulder, CO, 2003.

Kelly, Chris, "Value Monism, Richness, and Environmental Ethics," *Les ateliers de l'éthique/The Ethics Forum*, vol. 9, no. 2, 2014, pp. 100-129.

Leibniz, Gottfried Wilhelm, *Discourse on Metaphysics*, Manchester, Manchester University Press, 1953.

Lewis, David, "Attitudes *De Dicto* and *De Se*," *The Philosophical Review*, vol. 88, no. 4, 1979, pp. 513-543.

Nozick, Robert, *Philosophical Explanations*, Cambridge, MA, Harvard University Press, 1981.

Oddie, Graham, *Likeness to Truth*, Dordrecht, Reidel, 1986.

Oddie, Graham, *Value, Reality and Desire*, Oxford, Oxford University Press, 2005.

Olson, Jonas, "Fitting Attitude Analyses of Value and the Partiality Challenge," *Ethical Theory and Moral Practice*, vol. 12, no. 4, 2009, pp. 365-378.

Rabinowicz, Wlodek, "Value relations," *Theoria*, vol. 74, no. 1, 2008, pp. 18-49.

Rabinowicz, Wlodek, "Value, Fitting Attitude Account of," in Hugh LaFollette (ed.), *The International Encyclopedia of Ethics*, Oxford, Wiley-Blackwell, 2013.

Rabinowicz, Wlodek and Toni Rønnow-Rasmussen, "The Strike of the Demon: on Fitting Pro-attitudes and Value," *Ethics*, vol. 114, no. 3, 2004, pp. 391-423.

Tappolet, Christine, "Values and Emotions: Neo-Sentimentalism's Prospects," in Carla Bagnoli (ed.), *Morality and the Emotions*, Oxford, Oxford University Press, 2011.

Wedgwood, Ralph, "The "Good" and the "Right" Revisited," *Philosophical Perspectives*, vol. 23, no. 1, 2009, pp. 499-519.

Zimmerman, Michael J., *The Nature of Intrinsic Value*, Lanham, Rowman and Littlefield, 2001.

Zimmerman, Michael J., "Intrinsic vs. Extrinsic Value," Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Winter 2010 Edition, URL = <<http://plato.stanford.edu/archives/win2010/entries/value-intrinsic-extrinsic/>>.