

The Truth in ‘The Truth in Particularism’

Debbie Roberts

Draft: please do not cite without permission.

If one decision is fair and another unfair, that can't be the only difference between them. This is true. But observations of this sort are typically a gateway to more general claims: ‘There can't be an evaluative difference between two things (situations, decisions, legal systems, laws, actions, characters etc.) without there being some *non-evaluative* difference between them’. ‘The evaluative supervenes on the non-evaluative. It does so because it is grounded in the non-evaluative’. These general claims are very widely accepted, some hold them to be a priori conceptual truths, and they are put to important dialectical use.

Whilst almost all metaethicists think that something in the vicinity of these supervenience and grounding theses just must be the case, it is worth noting that it is not at all clear that there are theses here that all metaethicists agree on. For one thing, supervenience claims can be formulated in a variety of ways. To give a non-exhaustive list: They can be global, regional, or local, i.e., a claim about entire worlds, a claim about a space-time region, or a claim about particular individuals. They can be broad or narrow, i.e., include the supervening class in the base or not. They can be weak or strong, holding within worlds or across possible worlds. They can be ascriptive or ontological, i.e., a claim about a relation between judgments or a claim about a relation between properties. Moreover, what the evaluative is held to supervene and depend upon varies e.g., the non-evaluative, descriptive, natural, and factual. And, as Nicholas Sturgeon has shown, these are not equivalent.¹

There is some tension here between the widespread conviction that something in the vicinity of the supervenience and grounding theses above just must be the case, hence their role as dialectical weapons, and the fact that they are difficult to formulate. There are different ways that we might seek to resolve this tension. Sturgeon, for example, argues that there are no non-parochial versions of the supervenience claim by which he means that there is no version of supervenience which all established metaethical views should assent to. The version that comes closest, he argues, to garnering universal assent is the claim that the normative supervenes on the factual understood as a broad thesis, i.e., including the supervening class in the base. If one's metaethical view allows for evaluative facts this thesis is then trivially true.² This version of the thesis makes no claims about ontological dependence. And the view that what explains supervenience is that the normative is grounded in the base thus gets no purchase.

Another way one might seek to resolve the tension would be to challenge the intuition that something in the vicinity of the supervenience and grounding claims *just must* be the case. This is the route taken by Joseph Raz in ‘The Truth in Particularism’ (1999). One section of that paper is devoted to explaining the irrelevance of supervenience to the debate between generalists and particularists. Raz argues that we have no reason to believe in the supervenience thesis because, ultimately, it has no explanatory power.

¹ Sturgeon (2009).

² The thesis is also trivially true for error theorists about the evaluative: if there are no evaluative facts then, trivially, there can be no evaluative difference without a non-evaluative difference for there can be no evaluative difference.

My aim in this paper is to defend this argument. I proceed as follows: in §1, I reconstruct Raz's argument. This reconstruction takes some doing for it is not transparent just how the argument is meant to go. That this argument is hardly ever discussed is some evidence for this. One person who does (briefly) discuss Raz's argument is Tristram McPherson.³ McPherson claims that Raz fails to consider an important explanatory case for supervenience, that the general supervenience thesis explains a whole host of specific supervenience theses. The kinds of specific theses that McPherson has in mind will be familiar from the literature on supervenience. They are the kinds of theses that are invoked in the conceivability argument for supervenience. Arguments for supervenience are thin on the ground, but this is plausibly the most powerful one. Raz does not consider these kinds of theses, or the conceivability argument. Indeed, he asserts at one point, perhaps surprisingly, that no specific theses are available. I discuss McPherson's objection in §2. Another surprising claim that Raz makes, in an important early step in his argument, is that everyone agrees that we have no reason to think that evaluative predicates or concepts supervene on non-evaluative predicates or concepts. It is not at all obvious that everyone agrees with this. After all, supervenience is widely held to operate as a conceptual constraint on evaluative judgment. Many hold that ascriptive supervenience is the case: that supervenience is a constraint that we are committed to by some deep feature of our evaluative discourse. I discuss this view in §3.

Raz claims that we have no reason to believe in supervenience because it is not explanatory. But this, one might think, gets things the wrong way round. Supervenience, many hold, is a fixed point, either because they think it is a conceptual truth, and a conceptual constraint on evaluative judgment or because they think it is a fundamental metaphysical truth. It may not itself explain, but it needs explanation. And what explains the supervenience of the evaluative on the non-evaluative is that the evaluative is grounded in the non-evaluative. The evaluative depends on the non-evaluative. In §4 I consider the grounding argument for supervenience.

The key move in Raz's argument concerns observations that he makes about ordinary evaluative practice. As they stand, these observations are suggestive but brief. They need spelling out in more detail. I argue in §5 that these observations supplemented by considerations about thick evaluation give Raz the resources to respond to all three of the cases for supervenience (and grounding) canvassed in the previous sections.

1. Raz Against Supervenience.⁴

Raz begins his paper on particularism by pointing to two considerations he claims that any account of reasons or of practical deliberation must accommodate. The first of these, the evaluative and guiding functions of reasons and of evaluative properties, is not relevant here. The second consideration he terms the *intelligibility* of the domain of evaluative properties and of reasons. He glosses this as follows:

The intelligibility of the domain of value means that nothing in it is as it is just because that is how it is; there is nothing 'arbitrary' in the domain of value. It is, after all, the domain of reason(s). There is an *explanation* for everything, an explanation for why what is good is good, what is bad is bad, etc. The intelligibility of value exerts a strong pressure towards some sort of a generalist view. [I]t implies that regarding any two situations such that

³ McPherson (2022) §5.2.

⁴ Raz (1999) p. 219-225. My italics.

*some evaluative concept applies to one of them but not to the other there is some further difference between them which can be helpful in explaining why.*⁵

This last sentence echoes the supervenience and grounding intuitions that I began this paper with. Depending on how one interprets ‘any two situations’ it invokes some idea of necessary co-variance (no evaluative difference without some further difference) which is the core idea of supervenience. Moreover, it seems natural to move from this thought, that there must be some further difference that explains why the evaluative concept applies in one case but not the other, first to the thought that the appropriateness of the application of the concept in one case rather than the other *depends* in some way on this further difference and second that this further difference *cannot itself be evaluative* (for then the point applies again).

However, Raz moves from this point about intelligibility to the claim that the supervenience of the evaluative on the non-evaluative is irrelevant to it. He rejects, that is, the move from the claim that there must be some further difference that explains why the evaluative concept applies in one case but not the other, to the claim that *there can be* no evaluative difference without a non-evaluative difference.⁶ Why?

In short, Raz’s view is that the supervenience thesis does not explain *why* an evaluative concept applies in one case and not in another. Moreover, when we get clear on what the supervenience thesis is, or rather could at best be, we’ll find that it explains nothing and thus that we have no reason to believe in it.

It is worth noting that strictly speaking, supervenience does not include dependence. Strictly speaking, supervenience is a claim about necessary co-variance (no A-difference without a B-difference). But historically, many working in metaethics have understood supervenience to include dependence (The As depend on the Bs). Moore, for example, talks of moral properties as supervenient *or consequential* properties. In my view, Raz understands supervenience as necessary co-variance and as a dependence relation. This will become particularly relevant in the section on grounding. However, in this initial reconstruction of his argument, I will focus as far as possible only on necessary co-variance. As will become clear when we get to objections, it is helpful to clearly distinguish necessary co-variance and dependence.

Before we get to my reconstruction of Raz’s argument, it is worth making clear the precise supervenience claim that Raz is disputing. First, it is strong supervenience that he is concerned with, his interest is in the claim that supervenience is of the *essence* of the evaluative.⁷ Second, what he doubts, he says, is “the truth of *global* evaluative supervenience theses in general”.⁸ What he means by “global” is that these theses refer to the supervenience of all evaluative predicates, concepts, or properties on all the non-evaluative ones.⁹ This is not the typical understanding of ‘global’ in this context. Typically, a global supervenience thesis is one that concerns entire worlds rather than individuals or situations. Raz’s theses concern situations rather than worlds, and so would count as ‘local’ on the typical understanding

⁵ p. 220

⁶ He says just before he launches into the argument against supervenience is that he doubts “whether there is much we can learn from any general supervenience thesis we know of today”, p. 220.

⁷ p. 221 n. 8. He does say that even natural necessity would be enough to imbue the supervenience thesis with philosophical interest, but in general he has in mind the stronger thesis.

⁸ p. 221.

⁹ p. 221 n. 7.

In light of the above, his main claim is best construed as the claim that we have no good reason to believe that strong global (in his sense) supervenience is true, whether construed as a thesis about predicates, concepts, or properties. His argument for this claim is elusive. To orient the reader, I summarise it briefly first before explaining it in more detail.

He proceeds in four main steps.

First, he argues that we have no reason to believe that evaluative predicates or concepts supervene on non-evaluative predicates and concepts. The range of predicates, evaluative and non-evaluative, available in a language is a contingent historical matter. The same is true of concepts. If there happens to be covariance between evaluative and non-evaluative predicates/concepts for some natural language/some speakers at some point in time that would be a surprising historical accident. One might think no! Predicate/concept supervenience must always be the case! But, Raz argues, an examination of ordinary evaluative practice reveals that there is nothing that would necessitate that co-variance.

Second, Raz claims that arguments that would establish *necessary* co-variance of evaluative and non-evaluative predicates/concepts would only establish supervenience under ideal conditions, not our current conditions with our current predicates and concepts. These are arguments from considerations of objectivity, and considerations regarding complete understanding of evaluative concepts.

Third, he argues that supervenience under ideal conditions, concerning relations between predicates and concepts that we don't have access to (and, he thinks, likely could not have access to) can be of no help in explaining the current intelligibility of the evaluative domain.

Lastly, Raz turns to a metaphysical version of the supervenience thesis. Perhaps the necessary co-variance is to be found in relations between evaluative and non-evaluative *properties* rather than predicates and concepts. However, he argues, given the fact that criteria for property identification are tied to particular conceptual schemes, we don't have any reason to believe that supervenience holds for the properties that we currently know of. Perhaps it holds, again, under some ideal conditions, but potentially we have no epistemic access to that situation. A supervenience relation between properties that we don't know about can't help explain the intelligibility of the evaluative domain. What we are in explanatory need of are *specific* supervenience theses which tells us which particular evaluative properties supervene on which particular non-evaluative properties and why. But the general supervenience thesis doesn't entail these and moreover it doesn't even guarantee that they are possible – for it may be inaccessible to us. We should conclude that neither the predicate/concept nor the property general supervenience thesis explains anything. We thus have no reason to believe in evaluative supervenience.

- i. *Step 1: There is no reason to believe that evaluative predicates and concepts (of any natural language/speakers) supervene on non-evaluative predicates and concepts.*

The supervenience thesis applied to predicates is that evaluative predicates supervene on non-evaluative predicates if and only if, necessarily for any two situations, if there is an evaluative predicate which applies to one and not to the other then there is a non-evaluative predicate which applies to one and not to the other. Applied to concepts it is that if and only if, necessarily for any two situations, if there is an evaluative concept which applies to one situation and not to the other, then there is a non-evaluative concept which applies to one and not to the other.

Raz claims that it is generally agreed that there is no reason to think that the evaluative predicates of English (or of any other natural language) supervene on its non-evaluative predicates. Nor, he says, is there any reason to think that the evaluative concepts available to English speakers (or to members of any other group) today supervene on non-evaluative concepts available to them. Why might this be? Because which predicates and concepts belong to a language or cultural group at any particular time *is a contingent matter*. So, it follows that if predicate or concept supervenience obtains at any particular time, that too would also be a (surprising) contingent matter.¹⁰

Crucially, Raz claims, supervenience that obtains merely by historical accident could not have any philosophical importance. For predicate/concept supervenience to be philosophically important there must be something that makes it necessary that it should obtain. Raz considers, and rejects, an argument for the claim that there is some mechanism or factor that makes it necessary that predicate/concept supervenience should obtain for any natural language/speakers at any point in time. This argument concerns what is involved in the acquisition and mastery of evaluative predicates and concepts. For one thing that would make it the case that evaluative predicate/concept supervenience necessarily obtains would be if we could not acquire or master evaluative predicates and concepts without proceeding via acquisition and mastery of non-evaluative predicates and concepts. The idea here would be, roughly, that the ability to apply evaluative predicates correctly requires or consists in first establishing which non-evaluative one's apply and then concluding whether the evaluative ones apply.

But, says Raz, this is manifestly not the case. If there is evaluative predicate/concept supervenience this is not because our understanding of these predicates/concept and our ability to apply them correctly “presupposes knowledge of how they or the contexts of their application can be characterised or identified by non-evaluative predicates and concepts”.¹¹ For one thing, we don't acquire evaluative predicates or their associated concepts in this way, and for another this isn't how mature evaluative judgment works. Close attention to ordinary evaluative practice reveals that we learn how to apply evaluative predicates by example and by their association with other evaluative predicates and concepts. Indeed, typically, we identify contexts in which an evaluative predicate applies *in evaluative terms* and our ability to do this “far outstrips” our ability to do so in non-evaluative terms.¹² As an illustration, consider the following: why did you think that what she did to her was *cruel*? Well, because it involved *gratuitously harming* an *innocent*. Why was it gratuitous? There was *no justification* for it. What made it a harm? Well, it was a *malicious deceit*, and it negatively impacted her *welfare*. Why was she innocent? She hadn't done anything *wrong*!¹³ Moreover, often, says Raz, there is no way to determine the application of an evaluative concept except by reference to other evaluative concepts.¹⁴

At this point, Raz takes himself to have established that we have no good reason to think that predicate/concept supervenience necessarily obtains for any actual natural language/speakers at any point in time.

¹⁰ If supervenience is *necessary* co-variance, then strictly speaking, if there is no necessity then there is no supervenience. But I take it that what Raz means here by contingent supervenience is that there just happens to be co-variance (and dependence – he seems to understand supervenience as including a dependence claim. I come back to this below).

¹¹ p. 222.

¹² p. 221.

¹³ Cf. Roberts (2017) p. 209.

¹⁴ This last point is consistent with supervenience, as Raz, referencing McDowell, points out. p. 221, n. 11.

- ii. *Step 2: At best, arguments for the necessity of evaluative predicate/concept supervenience establish such supervenience under ideal conditions.*

Raz considers two further arguments for the necessity of evaluative predicate/concept supervenience. The first concerns objectivity and the second the relation between predicates/concepts and properties.

The first goes as follows:

Evaluative discourse is objective and admits of the possibility of mistakes, and of criteria for correctness. But, ultimately, only naturalistic discourse can be objective.¹⁵ Any other discourse can only be objective to the extent that it depends on naturalistic discourse. So, evaluative predicates and concepts must supervene on non-evaluative ones.

This version of the objectivity argument fails, Raz claims, for the same reasons that the argument examined in *i* above failed, i.e., close attention to actual evaluative practice reveals no such dependence on naturalistic discourse. If this objectivity argument is to succeed, it thus needs to be supplemented with more stringent conditions on objectivity. Two candidate such conditions are first, that without predicate/concept supervenience there will be no guarantee of convergence of inquirers under ideal conditions and second that the best explanation of beliefs in evaluative propositions can only be the existence of evaluative facts if predicate/concept supervenience is the case.

Raz himself rejects these further conditions on objectivity, however, even if we accept one or both, he argues, the objectivity argument could not establish the necessity of evaluative predicate/concept supervenience for languages (speakers) as we have them. *At best, it could establish supervenience in some ideal or extended language.* Raz does not spell out his reasoning here, he merely notes that the objectivity argument does not deal with the problems for the supervenience thesis (as applied to natural languages) raised by the contingency of the range of available predicates and concepts.

In the case of the convergence condition, the thought seems to be something like the following: we don't, Raz thinks, have any reason to think that evaluative predicate/concept supervenience obtains now with the range of predicates/concepts we currently have. If it were to obtain under ideal conditions, then it would have to obtain in virtue of a different range of predicates/concepts. In the case of the best explanation condition, my best guess is that Raz here again has his points about actual evaluative practice in mind: our current explanations for our evaluative beliefs do not (typically) invoke dependence of those beliefs on non-evaluative or naturalistic beliefs, because the application of our current evaluative predicates/concepts do not in general depend on the application of non-evaluative predicates/concepts. So, if the *best* explanation of our evaluative beliefs does invoke non-evaluative beliefs, and thus non-evaluative predicates and concepts, then that best explanation will have to involve an ideal or extended language.

The second argument of the necessity of evaluative predicate/concept supervenience also, Raz claims, fails to establish such supervenience for languages as we currently have them and only establishes it for some ideal/extended language. This argument concerns properties, as well as predicates and concepts, and claims that a complete understanding of evaluative concepts

¹⁵ Raz switches without warning here from speaking of supervenience on the non-evaluative to supervenience on the natural. I take it from the context that he thinks nothing much hangs on the difference here. And it is certainly true that defenders of supervenience tend to use these terms interchangeably.

includes understanding how the evaluative properties they refer to supervene on non-evaluative properties. If sound, Raz notes, this would show that no-one currently has a complete understanding of many evaluative concepts. (Though he does not make this explicit, again his reason here would appear to be what is revealed by a close consideration of actual evaluative practice.) Thus, if sound, this argument only establishes supervenience under ideal conditions.

- iii. Step 3: We have no useful or stable way of working out what these ideal conditions are, and in any case a thesis that holds only under as yet undiscovered ideal conditions is irrelevant to explaining the current intelligibility of the evaluative domain.*

Since we are talking of evaluative predicate/concept supervenience, one might think that ideal conditions could be secured by enriching our current language with additional non-evaluative predicates, so that it turns out that the evaluative predicates/concepts supervene on the non-evaluative ones. But Raz thinks this is not a promising strategy. He gives one reason for thinking (though implies he has others too):¹⁶

For one thing, it seems compatible with the reverse thesis. At any given time, the counter-thesis goes, if the evaluative predicates then in English supervene on its non-evaluative predicates it is possible to enrich English with additional evaluative predicates so that its evaluative predicates will no longer supervene on the non-evaluative ones.

At this point, it is useful to take stock. We began with the thought that the evaluative domain is intelligible. If one decision is fair, and another unfair, there must be some further difference between the two decisions which helps explain the difference. If we think of what needs explaining as why the predicate 'fair' correctly applies in one case and not in the other, one explanation might be that there is, necessarily, some non-evaluative predicate (s) that applies in the one case and not in the other, because evaluative predicate supervenience is true. We can run the same thought using concepts instead of predicates.

But, Raz has argued, what predicates (concepts) we have is an entirely contingent matter. Unless there is some mechanism which would necessitate predicate/concept supervenience, if it did happen to obtain, that would be a philosophically uninteresting historical accident. Close attention to evaluative practice reveals there is no such mechanism which guarantees such supervenience for languages as we have them. At best, if there is predicate/concept supervenience then it holds for some ideal language which we don't have, and it is not clear how we could get it. Most importantly, though, from a theoretical perspective, it is not clear why we need such a thesis. That is, Raz's thought seems to be: why think that it is this thesis (that holds, if it does, for some ideal likely inaccessible language) that provides the explanations required given the intelligibility of the evaluative domain? What reasons do we have to believe this?

- iv. Step 4: Moving to a claim about the necessity of a metaphysical evaluative supervenience thesis concerning properties is of no help; this thesis is explanatorily vacuous.*

One natural response at this point might be to turn to metaphysics. Indeed, someone of this mind might think we should have been talking about metaphysics all along, and that Raz's starting point, i.e., considerations to do with predicates, concepts, and the intelligibility of the evaluative domain, entirely misses the point. The supervenience of the evaluative on the non-evaluative obtains, this thought goes, because of the nature of what can exist, metaphysically

¹⁶ p. 223

speaking. Really, the supervenience thesis we should be considering is the following: a difference in evaluative properties between two situations *requires* a difference in non-evaluative properties—or, equivalently, exact similarity with respect to non-evaluative properties *guarantees* exact similarity with respect to evaluative properties. This has nothing to do with predicates or concepts, currently available or not.

Raz anticipates this worry, and his response is to take us back to our stock on concepts, actual or possible, and back to considering what reasons we must believe in supervenience, now applied to the metaphysical version of the thesis.

His first point is that the criteria for identification of properties depends, albeit indirectly, on our concepts. For example, suppose properties are identified by reference to causal powers. Since, Raz says, different conceptual schemes offer different ways of individuating the causal powers of things, different conceptual schemes offer different criteria for the individuation of properties. It follows that the metaphysical supervenience claim must either be tied to some particular way(s) of individuating properties or be held to apply regardless of how properties are individuated.

Which way do we go? The trouble at this point, says Raz, is that there's no reason to go either way, for neither will give us what we need. What we need, he says, are more specific explanatory theses: "We are in need of some understanding of which non-evaluative properties specific evaluative properties depend on and why".¹⁷ The general thesis does not give us these specific theses:

We seek to understand how evaluative events (i.e., events in which evaluative properties figure, like murder, or the performance of kind acts) fit in our understanding of causal explanations of events, and our tendency to regard causal explanations as the primary mode of explanations of events. The puzzles about the way evaluative facts and events fit into our scientific world-view are not solved by one global thesis of the evaluative but by specific theses which are not directly entailed by it.¹⁸

Moreover, the global thesis does not even guarantee the possibility of the more specific theses, says Raz. Given that individuation of properties depends on conceptual repertoire, we have no reason to think that, for the properties that we know of, that the evaluative ones supervene on the non-evaluative ones. The problem is analogous to that with predicates and concepts. And, assuming that the number of properties is infinite, if we could expand the set of properties by adding non-evaluative ones to secure supervenience, it may also be possible to further expand the set by adding more evaluative properties such that supervenience once again does not obtain. We should conclude from this that it is possible that if the global metaphysical supervenience thesis is true, it is true in virtue of properties we could not in principle come to know. Therefore, it could not figure in any explanation of properties that we do know. Raz ends his discussion of supervenience as follows:

These remarks show how little the general thesis gives us. It does not have much, if any, explanatory power. But the less explanatory power it has the less reason we have to believe in it. If some form of evaluative supervenience obtains, then for it to serve a significant explanatory role, we need specific rather than merely global evaluative supervenience thesis, but no such theses are available.¹⁹

¹⁷ p. 224.

¹⁸ Ibid.

¹⁹ p. 225.

2. Supervenience against Raz

Perhaps surprisingly, metaethicists typically think that there are an infinite number of specific supervenience theses forthcoming. Take Hare's example of St Francis. It is impossible, says Hare, that two individuals could be alike in every respect, placed in exactly the same circumstances and behaving in exactly the same way and differ only in that one is good and the other not. Many are happy to conclude from this that there can be no evaluative difference without a non-evaluative difference. We can multiply such examples at will. Here is one from Tristram McPherson:

Suppose that a bank manager wrongfully embezzles their client's money. If we imagine holding fixed how much the bank manager stole, and how; the trust their customers placed in them; what they did with the money; all of the short- and long-term consequences of their actions; and so on, it seems that there could not be a second action that perfectly resembled this embezzlement, except that the second action was right rather than wrong. Cases like this one seem to show a *necessary* connection: they suggest that the ethical character of the bank manager's act cannot vary without some other facts varying as well.

McPherson objects to Raz's claims about the explanatory poverty of the general supervenience claim.²⁰ He argues that Raz has got things backwards. We shouldn't be expecting supervenience either to provide more specific theses that allow us to concretely explain the evaluative features of reality, or to guarantee that we can find such explanatory theses. Instead, *we start with the specific theses*. We note that can identify specific supervenience theses for *any* evaluation: the evaluative quality of a character, action, situation and so on *cannot* vary without something else varying as well. We have an infinite list of such necessary connections. These facts may not themselves be explanatory, but they do seem difficult to intelligibly deny. The general supervenience thesis *unifies* these all of these into a single pattern. The general thesis, that is, serves to explain all the specific supervenience facts. Moreover, this general thesis is itself independently plausible. It reflects the fact, and McPherson here points to inspiration from Sidgwick, that while we can admit brute contingency into our account of the physical world, we cannot admit similar contingency into our account of the evaluative. It cannot be the case that a certain evaluative property just happens to covary with the base properties instantiated around here.²¹

I think Raz would agree with the above. Recall Raz's remarks about the intelligibility of the evaluative domain.

The intelligibility of the domain of value means that nothing in it is as it is just because that is how it is; there is nothing 'arbitrary' in the domain of value. It is, after all, the domain of reason(s). *There is an explanation for everything*, an explanation for why what is good is good, what is bad is bad, etc. The intelligibility of value exerts a strong pressure towards some sort of a generalist view. [I]t implies that regarding any two situations such that some evaluative concept applies to one of them but not to the other there is some further difference between them which can be helpful in explaining why.²²

²⁰ McPherson (2022) §5.2 & §2. Also, McPherson (2012, p. 211).

²¹ 'Base' is McPherson's term.

²² p. 220

I think it is clear from this that Raz accepts that the evaluative does not ‘float free’. That is, something can’t just be good, or unfair. It must have some other qualities. True, Raz’s remarks here go further than just noting covariance, in fact they suggest to me that he is thinking of dependence or grounding. But part of that is the recognition that if two situations differ in some specific evaluative respect *there must be some further difference* between them.

Note that the specific theses that McPherson has in mind can all be characterised in this way. If in one situation we have wrongful embezzlement, and in another situation we don’t, there must be some further difference between them. If we have a genocide in one case and not in another, there must be some further difference between them. And so on. These theses are not explanatory, as McPherson admits. They are indeed difficult to intelligibly deny, but Raz does not deny them. And nor, I think, would he deny the general thesis that unifies them. He would also not deny that that general thesis reflects a ban on brute contingency.

Raz would take issue, however, with the explanatory power that this general thesis has. I suspect that he would think that though this all fine as far as it goes, it does not go very far at all. It is merely a version of what he begins with, i.e., the intelligibility of the evaluative domain.²³ Recall that the specific theses Raz is after are not merely about co-variance: “We are in need of some understanding of which non-evaluative properties specific evaluative properties depend on and why”. “[R]egarding any two situations such that some evaluative concept applies to one of them but not to the other there is some further difference between them which can be helpful in explaining why.” Raz is after specific thesis which tell us what the further difference is and the explanation of why that difference makes the difference it does. From this perspective, McPherson’s general thesis gives us very little. It is a partial statement of the situation, of the thing that Raz thinks needs explaining. Partial because it addresses only necessary co-variance and not dependence.

At this point, one might object that I have left out an important element of the McPherson-type approach. Most metaethicists would claim that the kind of conceivability argument that McPherson employs to generate specific supervenience thesis shows that it is inconceivable that two situations could differ in their evaluative properties without also differing in their *non-evaluative* properties. Or, equivalently, that if we have two situations or individuals exactly alike in all their *non-evaluative* properties then they must be exactly alike in their evaluative properties too.²⁴ We notice these striking modal connections. The general thesis that explains (by unifying) them is that there can be no evaluative difference without a non-evaluative difference.

We can put this conceivability argument as follows:

P1: If we conceive of another individual identical to X in terms of non-evaluative properties then we conceive of an individual identical to X in terms of evaluative properties. We cannot conceive of another individual identical to X in respect of all non-evaluative properties but different from X in respect of evaluative properties; we know a priori that there cannot be such an individual.

P2: If we conceive of an individual different from X in terms of evaluative properties then we conceive of an individual different from X in terms of non-evaluative properties. We cannot

²³ A weaker version, for Raz’s characterisation suggests dependence, or grounding, as well as necessary co-variance.

²⁴ There are issues with how to characterise the base Cf. Sturgeon (2009). I’m ignoring these for the moment.

conceive of X being different in respect of evaluative properties without a difference in respect of non-evaluative properties.

P3: what explains (unifies) all specific supervenience theses generated by particular substitutions for X is a general supervenience thesis: whenever something has an evaluative property, it has a non-evaluative property or collection of non-evaluative properties that necessitate(s) the evaluative one.

C: Evaluative properties supervene on non-evaluative properties.²⁵

Now this version of the general thesis, I think Raz would not accept. His final verdict on this thesis is that it exhibits extreme explanatory poverty. But this is not because it would amount, from Raz's point of view, merely to a (partial) description of the situation we find ourselves in i.e., a description of what needs explaining. This was true of the general thesis that if two situations differ in some specific evaluative respect *there must be some further difference* between them. I think Raz doubts that we have good reason even to accept 'there can be no evaluative difference without a non-evaluative difference' as a correct description of the situation we find ourselves in. He may well accept that 'there can be no evaluative difference without a *base* difference'; i.e., I suspect he would be happy to agree that evaluative properties cannot vary without some underlying properties varying. But his argument above is meant to show that we don't have good reason to accept the characterisation of the base as non-evaluative.

His point is that in our current epistemic situation, with the evaluative predicates and concepts that we have and with the properties that we currently know of, attention to ordinary evaluative practice does not show that evaluative predicates/concepts/properties covary with non-evaluative predicates/concepts/properties. Furthermore, there is no mechanism or factor, he argues, which makes such co-variance necessary. So, if it is true (and we've got no reason to think it is), it is true in virtue of predicates/concepts we don't have and properties we don't know about.

Where does this leave us? I think it is helpful at this point to leave aside explicit focus on explanatory concerns. Instead, we should focus on what I am going to call the conceivability argument for supervenience. The objection to Raz's argument now is that regardless of why necessary co-variance is the case, the fact that there is such co-variance is incontrovertible: we cannot conceive of it being otherwise!²⁶ For it is inconceivable that two situations could differ in their evaluative properties without also differing in their non-evaluative properties. The challenge at this point can be put thus: give us an example, Raz, of two cases which differ evaluatively but do not differ non-evaluatively.

I come back to this challenge, and what I think Raz should say about it, in §5. Now I turn to a further reason one might have for thinking that the evaluative supervenes on the non-evaluative, one that takes issue with the very first step in Raz's argument .

3. Supervenience as a Conceptual Constraint

²⁵ Cf. Roberts (2018) p. 17.

²⁶ I'm here also leaving aside more general worries about how we know modal truths.

Raz's claim at the beginning of his argument that everyone agrees that there is no predicate/concept supervenience is surprising is because there is a very well-entrenched view in metaethics that would deny this.

Supervenience, according to those who endorse this view, is a conceptual constraint on our evaluative judgments to which we are committed by some deep feature of our evaluative discourse. We see this view in the writings of, for example, R. M. Hare, Simon Blackburn, and Michael Smith. Those who flout the supervenience requirement when they make evaluative judgements are supposed thereby to reveal themselves to be incompetent. On this view, we cannot judge objects to displaying evaluative differences without supposing them to have non-evaluative differences as well. If you were to judge that two situations were non-evaluatively identical but evaluatively different, you would be making a conceptual error, revealing yourself to be incompetent with evaluative predicates and concepts. This is to understand supervenience ascriptively, a piece of terminology which originates with James Klagge, as a relation between types of judgments rather than ontologically as a relation between types of properties.²⁷

Take two situations. According to ascriptive supervenience, logically speaking, a person's evaluative judgments (i.e., the evaluative predicates and concepts they take to apply) about the situations cannot differ unless their non-evaluative judgments (i.e., the non-evaluative predicates and concepts they take to apply) differ. And, equivalently, if someone takes all the same non-evaluative predicates and concepts to apply in two cases then they must take all the same evaluative predicates and concepts to apply. It is clear that those who accept ascriptive supervenience take it to hold for our current evaluative predicates and concepts.

Raz straightforwardly denies ascriptive supervenience. Recall his point about the contingency of the range of evaluative and non-evaluative predicates and concepts that a language or group of speakers has available to them. That point, coupled with his observations about evaluative practice is supposed to scupper the prospects for ascriptive supervenience.

Raz's observations about ordinary evaluative practice are supposed to show that as a matter of fact our evaluative judgments do not covary with non-evaluative judgments. Recall those observations: Thinking about ordinary evaluative practice shows that our understanding of evaluative predicates/concepts and our ability to apply them correctly does not presuppose knowledge of how they or the contexts of their application can be characterised or identified by non-evaluative predicates and concepts. For one thing, we don't acquire evaluative predicates or concepts in this way, and for another this isn't how mature evaluative judgment works. Close attention to ordinary evaluative practice reveals that we learn how to apply evaluative predicates by example and by their association with other evaluative predicates and concepts. Indeed, typically, we identify contexts in which an evaluative predicate applies *in evaluative terms* and our ability to do this "far outstrips" our ability to do so in non-evaluative terms. Moreover, often, says Raz, there is no way to determine the application of an evaluative concept except by reference to other evaluative concepts. Now this is all consistent with the possibility of predicate/concept supervenience. But if it holds, it doesn't hold for our current predicates and concepts but for some other set we don't currently have access to.

Here it seems we have a straightforward stand-off. On the one side there are those, like Hare and Blackburn, who think our evaluative predicates and concepts just do supervene on non-evaluative ones and on the other there is Raz who does not. Curiously, both these opposing views seemed to be in some sense based on observations about ordinary evaluative practice. One

²⁷ See Klagge (1988).

way to break this stand-off might be to consider a further argument for supervenience, Michael Smith's grounding argument.

4. Evaluative claims cannot be brutally true

Smith begins this argument by pointing out that it is simply incoherent to suppose that evaluative claims could be *barely* true. Evaluative claims must always be *made true* by other claims. An evaluative claim, about a life being good or an act being wrong for example might be made true by more specific evaluative claims, e.g., it was a life of accomplishment or a cruel act, but these more specific evaluative claims cannot be brutally true either. We can go on to ask what makes them true. Ultimately, argues Smith evaluative claims must be made true by non-evaluative claims, claims about non-evaluative features that make the evaluative claims true.²⁸

Though Smith does not make this explicit, it is a clear implication of his discussion that he is understanding the 'making true' relation as a necessitation relation. We can extract the following argument from his discussion:

P4: It is a conceptual constraint on evaluative judgments that evaluative claims can't be barely true; they must always be made true by other claims, ultimately by claims about non-evaluative properties that make the evaluative claim appropriate.

P5: (From P4) if an individual has an evaluative property, it has that evaluative property, ultimately in virtue of non-evaluative properties that it has.

P4: In a possible world that agrees in the truth of all the same claims about non-evaluative properties, the same evaluative claim will be true.

P5: (From P4) in a possible world where another individual has all the same non-evaluative properties, it will have the same evaluative property.

C: Whenever something has an evaluative property, it has a non-evaluative property or collection of non-evaluative properties that necessitates the evaluative one.²⁹

In brief, supervenience is true, according to Smith because the evaluative is grounded in the non-evaluative. Smith claims this is a conceptual truth. Like Raz, he supports his argument with an appeal to ordinary evaluative practice:

Note that the fact that evaluative claims cannot be barely true is reflected in ordinary evaluative practice. Suppose I say that a particular life is good, but then look totally flummoxed when asked to provide the features of the life that make it good. Perhaps I say, "You clearly don't understand. It isn't made good by *other* features. *It's just good!*" If I am using "good" as an evaluative term, then I would plainly violate the rules that govern the use of the word "good." When I say of a life that it is good, using "good" as an evaluative term, I thereby incur an obligation to say why my ascription of goodness to

²⁸ Smith (2004) p. 225-229. Smith talks of the natural rather than the non-evaluative.

²⁹ Cf. Roberts (2018) p. 12-13.

that life is appropriate in the light of the features that the life possesses. If the life is good, these are the features of the life that make it good.³⁰

But while it is plausible that ordinary practice establishes that normative claims can't be barely true, it is not clear that it reveals that normative claims must be made true, ultimately, by claims about non-evaluative properties. It is precisely this claim that Raz would dispute. In fact, according to Raz, ordinary evaluative practice reveals that we stop, most of the time, with further evaluative claims.

Smith's response to this kind of move is to insist that we can always ask, for any more specific evaluation *v* 'but what makes it true that it is *v*?' Why did you think that what she did to her was *cruel*? Well, because it involved *gratuitously harming* an *innocent*. Why was it *gratuitous*? There was no *justification* for it. What made it a *harm*? Well, it was a *malicious deceit*, and it negatively impacted her *welfare*. Why was she *innocent*? She hadn't done anything *wrong*! At this point Smith says that if we follow this line of questioning to its logical conclusion we will have to admit that eventually this process will of necessity bottom out in claims about non-evaluative features because ultimately what makes it true that what she did was cruel was that she *acted* in some particular way, and acts, says Smith are at bottom bodily movements with certain characteristic causes – desires, beliefs, thoughts.³¹

5. Ordinary Evaluative Practice and the Thick

In each of the previous three sections, we have seen that the key move in Raz's argument, which is also the locus of disagreement for all three objections, concerns his appeal to ordinary evaluative practice. We thus need to consider this more closely. Here are the key claims: Typically, says Raz, we identify contexts in which an evaluative predicate applies in evaluative terms and our ability to do this "far outstrips" our ability to do so in non-evaluative terms. Moreover, often, there is no way to determine the application of an evaluative concept except by reference to other evaluative concepts.

If Raz is correct about this, what follows? First, it follows that Smith's claim, that evaluative claims must ultimately be made true by claims about non-evaluative features, is dubitable. At least, if Raz is correct, then often we are not able to articulate what it is that ultimately makes our evaluative claim true. What reason then do we have for thinking that the process *must* bottom out in this way? As far as Raz is concerned, we don't have one – the rest of his argument is supposed to show this. Second, if Raz is correct about ordinary evaluative practice then, as noted above, ascriptive supervenience is false – as far as our current predicates and concepts are concerned, evaluative predicates and concepts do not supervene on non-evaluative predicates and concepts. And third, I will argue, if Raz is correct about ordinary evaluative practice, the conceivability argument fails to show that he evaluative supervenes on the non-evaluative. But is Raz correct? Attention to the nature of thick evaluation gives us reason to think that he is.

Thin evaluative judgments are often made on the grounds of thick evaluative judgments. You judge an act to be wrong because cruel, or good because kind or courageous. Can we always spell

³⁰ Ibid p. 225.

³¹ Ibid p. 229. Smith is making certain assumptions about the nature of action here, ones that not all philosophers of action would accept. I think this is an important issue for a few reasons. One is that it is not clear how this controversial substantive view of the nature of action fits with his view that the supervenience thesis is a conceptual truth. Another is that there are alternative views of the nature of action where it is not possible to specify what an intentional action is in wholly non-normative terms.

out in non-evaluative terms what some act must be like to be cruel, or kind, or courageous? According to Raz, typically we cannot. Consider what an act must be like to be cruel. Well, it must involve the infliction of some kind of *prudential bad* or *harm*, but not any sort prudential bad will do (we wouldn't think of the infliction of a trivial prudential bad as cruel). It must be a prudential bad of the right sort of cruel-making sort. Moreover, this prudential bad must be inflicted with the right sort of intention or motivation – if you inflicted a prudential bad on someone by accident, that wouldn't count as cruel. But, of course, you can inflict a prudential bad out of indifference and that be cruel. Can we characterise what the motivation or intention or indifference must be like in fully non-evaluative terms? And the prudential bad?

It seems at least very difficult to do this. However, what is required for Smith's argument to get a grip, is not an account of the non-evaluative grounds of cruelty in general, i.e., applicable to all cases of cruelty. All that is required for Smith's argument to go through is that we bottom out in claims about non-evaluative features in every particular case of cruelty. In the end, it must be that what *necessitates* the evaluative property are non-evaluative properties. So let us consider the particular example of cruelty from above. Why did you think that what she did to her was cruel? Well, because it involved gratuitously harming an innocent. Why was it gratuitous? There was no justification for it. What made it a harm? Well, it was a malicious deceit, and it negatively impacted her welfare. Why was she innocent? She hadn't done anything wrong! Is it true that if we carry on this process, we will get down to non-evaluative properties which necessitate all the evaluative ones? What made the deceit malicious? Well, her intentions were hurtful and nasty. What made them hurtful and nasty? They were intentions to harm. It does not seem to me to be obvious that we must end up with a non-evaluative claim about bodily movements caused by desires and beliefs. Note, even if this is some part of the true account (and I do not deny that there will have been an action, but I do wonder if this is the correct account of action) why think that something like this is sufficient to *necessitate* the evaluative properties? Instead, I think that Raz is correct. Thinking about thick evaluation helps us to see that we identify contexts in which an evaluative predicate applies in evaluative terms and our ability to do this “far outstrips” our ability to do so in non-evaluative terms. Moreover, often, there is no way to determine the application of an evaluative concept except by reference to other evaluative concepts.

Other examples reinforce the point. A person gives a child some sweets. We describe fully all the bodily movements involved in non-evaluative terms (assuming for the moment that Smith's account of action is correct). What is the evaluative property of this action? It seems to me it that at this point it could be anything. It could be kind and good, it could be manipulative and bad. It could be sinister and evil. It could be thoughtless, it could be callous, it could be considerate. It could be sarcastic and cruel. It could be a joke. Now, one immediate objection might be: we need a fuller description of the context, and we can give that fully in non-evaluative terms, such that it is sufficient to necessitate the evaluative property. We need to know who this person is, what their relation is to the child. Is it a parent or grandparent, a teacher, or a stranger? We need to know some more about the child – are they hungry? Upset? How old are they? We need to know where and why this happening. Is it a school playground? The kitchen at home? An abandoned mine? Say it is a parent who gives a young child the sweets, they are in the kitchen at home. We don't yet have enough to determine the evaluative properties. Crucially, we need to know why this person gives the child the sweets. Say it's to keep the child quiet so the parent can complete some task in peace. Still, I don't think we have enough to necessitate the evaluative property (and I'm not sure that 'task' is non-evaluative – it connotes some piece of work that is in some sense worth doing). What if the child is diabetic? What if the thing the parent wants to do is trivial or self-indulgent and what the child really wants is a hug?

We can make this point by borrowing considerations from the generalism-particularism debate. No matter what non-evaluative description we give, it is always possible that some other feature of the context is relevant in a way that changes the evaluative features. But now, the objector will say, well then what we do is we expand the base to include *all* the non-evaluative features of the entire world in which this action is situated. That will necessitate the evaluative properties! It's not clear to me that we are licensed to make this assumption if we are sticking with the predicates and concepts currently available to us, and the properties that are tied to our current conceptual scheme, i.e., that we currently know of. For example, consider again the terms 'task', 'trivial' and 'self-indulgent' in the scenario above. It seems to me very difficult to spell out fully, in uncontroversially non-evaluative terms something in such a way that *necessitates* that it is trivial.

An objection at this point is likely to be that I (and perhaps Raz) am confusing different kinds of necessities. Really, this objection goes, the supervenience and grounding theses of interest are *metaphysical* theses. It is compatible with the claim that certain non-evaluative properties of an action ground, and thus metaphysically necessitate, its evaluative properties, that no non-evaluative description will imply, just as a matter of the predicates and concepts involved, any particular evaluative conclusion.³²

This is correct. Moreover, Raz accepts this. Strictly speaking his claim is not that the evaluative does not metaphysically supervene and depend on the non-evaluative. His point is that we don't have good reason to believe that it does. Note, first, that one way to put this objection is to say that even if Raz is correct that ascriptive supervenience is false – given our current predicates and concepts the evaluative ones don't supervene (or depend) on the non-evaluative ones – that doesn't license the conclusion that there is no metaphysical dependence or supervenience of the evaluative on the non-evaluative. Again, that's correct as far as it goes. But Raz's point is that it doesn't go that far. We need to ask what reason we have to believe in the metaphysical versions of supervenience (and dependence) if the ascriptive version is false. Here though we need to consider Raz's point that the difference between concepts and properties can be exaggerated. Raz focuses on how the criteria for identification of properties depends, however indirectly, on our conceptual repertoire. As far as the properties we know of are concerned, we have no guarantee that we can identify the precise non-evaluative properties that necessitate the evaluative properties in any particular case. So, what reason do we have to think that they *must* be there?

Thinking about thick evaluation gives us reason to think that Raz is correct when he says that we need some understanding of which non-evaluative properties specific evaluative properties depend on (and why), but we don't have this. No such specific supervenience (grounding) theses are available.³³

At this point, we've dealt with Smith's grounding argument for supervenience, and we've dealt with the Hare/Blackburn view that ascriptive supervenience is true. But what about the conceivability argument? Doesn't that give us a reason to think that the non-evaluative properties that necessitate the evaluative ones just must be there?

First, note that the *conclusion* of that argument is property supervenience. We shouldn't assume then in the premises of the argument that property supervenience holds. What is supposed to establish that conclusion is the claim that we can't *conceive* of two individuals non-evaluatively

³² Cf. Roberts (2018) p. 15.

³³ Raz (1999) p. 224-5

identical but evaluatively different, or evaluatively identical but non-evaluatively identical. But again, the point is that we are constrained in what we can conceive *by our current predicates and concepts*. Take another example used by McPherson:

[I]t seems impossible that another world might be identical to this one except that in the other world, a genocide otherwise identical to the actual Rwandan genocide differed solely in being ethically wonderful, rather than being an atrocity.³⁴

That seems correct. But that doesn't show that we are imagining two cases identical *only* in their non-evaluative features. Take out genocide, for starters. And take out brutality, cruelty, harm, and any evaluative characterisation of the content of intentions. Should we take out intentions as well? Can intentions be non-evaluatively characterised? It seems at least difficult to characterise them non-normatively.³⁵ What are we left with that we are *certain* is *wholly* non-evaluative? Maybe the killing of one group of people by another group of people. Now it is natural to think that this is an atrocity. But it could conceivably be instead a tragic accident or mistake.

We can make similar points about Hare's St. Francis example. Take St. Francis. It is not the case that we are tasked with imagining two individuals exactly identical in every respect except that one is a good person and the other not. We need to also take out all the virtue terms and concepts. It is perfectly plausible that in leaving *only* controversially non-evaluative descriptions we end up with something like 'a human being with some character traits'.³⁶ And if that is the case then the challenge to provide an example of two cases which differ evaluatively but do not differ non-evaluatively is quite easy to meet.³⁷

6. Conclusion

It is worth emphasising at this point that Raz does not deny that evaluative properties (or predicates, or concepts) are supervenience or dependent. He accepts that evaluative claims cannot be brutally true. He accepts that an individual cannot just be good and there be nothing that makes that individual good. What the questions is whether we have good reason to believe that what the evaluative supervenes and depends on is the *non-evaluative*.

There may be some residual worries to deal with, concerning consistency: like cases ought to be treated alike, and objectivity: only if the evaluative supervenes and is grounded in the non-evaluative can evaluative discourse be objective. It is true that like cases ought to be treated alike. Where the same reasons are in place the same evaluative judgments should be made. But the above discussion gives us reason to doubt that these reasons will always be non-evaluative. On the second point, I agree with Raz that these are not conditions for the objectivity of evaluative discourse.

³⁴ McPherson (2012) p. 211-12.

³⁵ I think it is safe to say that it is typical in discussions about supervenience of the evaluative on the non-evaluative to use 'evaluative' broadly to include the normative.

³⁶ 'Human being' rather than 'person' because it is not clear that 'person' is a non-evaluative, or non-normative, notion. One might balk at this and claim that it will always be possible to give the full description in non-evaluative terms, of how this being acts – we could give a full description of all the things this being does in non-evaluative terms. However, we've already seen reason to doubt that this necessarily must be possible above.

³⁷ One might object that I am assuming throughout that thick predicates and concepts are evaluative and that no reductive analyses of these are available. On the first point, this is true. On the second, I think the discussion and consideration of examples gives good a sense of the reasons for holding that reductive analyses of the thick are not possible. However, I defend both claims in ms.

References

- Klagge, James, 1988, "Supervenience: Ontological or Ascriptive", *Australasian Journal of Philosophy*, 66: 461–470.
- McPherson, Tristram, 2012 "Ethical Non-Naturalism and the Metaphysics of Supervenience", in Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics* (Volume 7), Oxford: Oxford University Press, pp. 205–234.
- , 2022, "Supervenience in Ethics", *The Stanford Encyclopedia of Philosophy* (Summer 2022 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2022/entries/supervenience-ethics/>>.
- Roberts, Debbie, 2017, "Depending on the Thick" *Proceedings of the Aristotelian Society* (Supplementary Volume), 91: 197–220.
- , 2018, "Why Believe in Normative Supervenience?", in Russ Shafer Landau (ed.), *Oxford Studies in Metaethics* (Volume 13), Oxford: Oxford University Press, pp. 1–24.
- , ms. "Thick Facts"
- Raz, Joseph, 1999, "The Truth in Particularism", in *Engaging Reason*, Oxford: Oxford University Press, pp. 48–78.
- Smith, Michael, 2004, "Does the Evaluative Supervene on the Natural?" in M. Smith, *Ethics and the A Priori*, Cambridge: Cambridge University Press, pp. 208–233.
- Sturgeon, Nicholas, 2009, "Doubts about the Supervenience of the Evaluative", in Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics* (Volume 4), Oxford: Oxford University Press, pp. 53–92.