Supervenience, Repeatability, & Expressivism*

Emad H. Atiq

Law School & Sage School of Philosophy, Cornell University

Abstract

Expressivists traditionally explain normative supervenience by saying it is a conceptual truth. I argue against this tradition in two steps. First, I show the modal claim that stands in need of explanation has been stated imprecisely. Classic arguments in metaethics for normative supervenience and those that rely on it as a premise presuppose a constraint on the supervenience base that is rarely (if ever) made explicit: the repeatability of the non-normative properties on which the normative supervenes. Non-normative properties are repeatable when it is possible for numerically distinct individuals to share them. Second, I show if the modal truth that stands in need of explanation entails that there are individuals exactly alike in repeatable non-normative respects that cannot normatively differ, then standard expressivist accounts of normative supervenience as a conceptual truth are unsuccessful. Expressivist metasemantics for normative terms, together with constitutive facts about the non-cognitive attitudes essentially involved in normative thought, strongly suggest that repeatable supervenience could not be a conceptual truth. I argue, finally, that although repeatable supervenience bears the marks of a conceptual truth, expressivists should be content to treat it as an ordinary normative truth, and to explain it the same way they explain other normative truths.

1. Introduction

Almost everyone writing in metaethics agrees that normative properties supervene on non-normative properties: individuals identical in their non-normative features cannot differ in their normative features.[1] However, this initial statement of the consensus turns out to be misleading. When we examine the standard arguments for supervenience along with the theoretical uses to which supervenience has been put in metaethics (including by expressivists, the principal protagonists of this story), it becomes apparent that theorists have had in mind a more particular modal claim—that normative properties supervene on *repeatable* non-normative properties.[2] Non-normative properties are repeatable when it is possible for numerically distinct individuals to share them.

As an expressivist, I am primarily interested in what this precisification means for expressivist accounts of normative supervenience. What it means is that the standard party line needs considerable refinement. Expressivist explanations of the modal claim as a conceptual truth do not withstand scrutiny. Expressivist metasemantics for normative terms, together with constitutive facts about the non-cognitive attitudes essentially involved in normative thought and talk, strongly suggest that repeatable supervenience could not be a conceptual truth.[3] In light of this, the most attractive option, as an expressivist, is to regard the relevant modal claim as an ordinary normative truth, rather than a conceptual one. I argue that the theoretical costs are minimal. The expressivist has unique resources for explaining why the claim bears all the hallmarks of a conceptual truth.

§1 offers three arguments for why the modal truth that has attracted wide interest in metaethics must be construed in terms of repeatability. §2 introduces the general challenge facing expressivists given repeatable supervenience: one can plausibly and without conceptual confusion have motivations directed at non-repeatable features, like the haecceitistic characteristics of individuals and situations. Such 'identity-dependent' motivations constitute intelligible repeatable supervenience flouting normative judgments. §3 explores whether the challenge can be avoided by conditioning the intelligibility of normative assertions on the expression of key intention-like states which appear constitutively constrained to preclude identity-dependence. Gibbard's 'plan-based' expressivism offers a helpful illustration of this strategy. The strategy fails, however, because it cannot be a necessary condition on the intelligibility of normative judgments that they be constituted by such intention-like states. §4 argues that it is unhelpful to simply stipulate that the relevant motivations are repeatability-constrained or appeal to the functions of normative language. §5 concludes with an alternative account of repeatable supervenience as an ordinary normative truth.

2. Clarifying the Explanandum

Normative properties—whether of acts, states of affairs, or whatever else—are said to supervene on the non-normative properties of their bearer. According to Michael Smith (2004, p. 8), the truth of supervenience "is accepted by nearly everyone writing on the nature of value" and its denial would be evidence of a very basic conceptual confusion.[4] Arguments for supervenience tend to be rare given the assumption of self-evidence. One argument, the Argument from Cases, appeals to what seems true in individual cases. Consider a wrongful act, say, an instance of lying. It seems inconceivable that there could be another lie exactly alike in all non-normative respects (it is told for the same purpose and in the same type of situation, has the same effects, and so on) but that fails to be wrong (Smith 2004: 225–9). Fixing the non-normative features of an act settles its normative features. Such cases reflect a more general truth about the character of the normative, or so the argument from cases maintains.

There are different ways of stating what this general truth amounts to, but here is a widely-accepted first approximation:

$$\text{SUPERVENIENCE} : \Box\,(\forall N \in \alpha)\,(\forall x)\,[Nx \rightarrow (\exists B \in \beta)\,(Bx\,\&\,\Box\,(\forall y)\,(By \rightarrow Ny))],$$

where α is the class of normative properties, β the class of non-normative or base properties, and $\Box$ is a necessity operator.

In English: necessarily, for any normative property, if that property holds of some x, then there is some non-normative property that x has such that, necessarily, for any y that has the same non-normative property, y also has the normative property.[5] The relevant non-normative property might be a complex property: it could be the property of *being B_1 or B_2 or ....*[6]

There is relative agreement that the outermost necessity operator should be interpreted as conceptual necessity (Blackburn 1984, Smith 2004, Enoch 2011, Dreier 2015). It is supposed to follow from the nature of the concepts involved that the connection between normative and non-normative properties holds. The inner necessity operator should be interpreted as metaphysical necessity. The conceptually necessary truth is that it is metaphysically impossible for two things to be non-normatively identical but normatively distinct.

My concern is with β: the class of non-normative properties on which the normative supervenes. One might be tempted to leave β unrestricted if one's goal were to state the supervenience thesis as uncontroversially and modestly as possible (McPherson 2012, Rosen forthcoming).[7] But leaving β unrestricted obscures the precise supervenience phenomenon that has historically been the target of metaethical analysis and explanation.

Unrestricted supervenience is true but trivial. If no restrictions on the class of non-normative properties are assumed, then among the base properties are haecceitistic properties that can only be instantiated by particular individuals, like the property of *being Fred*. There are also properties indexed to possible worlds, like the property of *being Fred in the actual world*. It is certainly true that two things cannot differ in normative respects without differing in some such non-normative respect. But that is because things alike in all of these respects are one and the same thing at a world. The property of *being Fred in the actual world* cannot be had by anyone other than Fred in the actual world, and, trivially, Fred in the actual world is either good or not good—he cannot be both. No one ever imagined that supervenience could be so easily explained.[8]

The constraint needed to block triviality cannot be too restrictive. The narrower the class of base properties, the stronger the corresponding supervenience claim. If the class is too narrow, the truth will not hold simply by virtue of the concepts involved (Blackburn 1993).[9] For example, it may be true that normative facts supervene on the physical facts. But this is not a conceptual truth, since it is conceptually possible that a non-physical thing, God, exists, and makes a difference to the normative truths (Sturgeon 2009).[10]

The minimal constraint which appears to have been widely presupposed concerns the repeatability of the base properties. There are echoes of the repeatability

condition in Hare's (1952) original discussion of ethical supervenience.[11] But we can afford to be more explicit. For any B $\in$ β, B is repeatable when, necessarily, if <x is B> then possibly there is a y distinct from x, such that <y is B>. The class of repeatable properties excludes *being Fred* and haecceities of situations or possible worlds, because these cannot be had by distinct individuals. Other properties that aren't repeatable include exhaustive qualitative non-normative profiles of individuals of the sort that cannot be shared between distinct individuals, for identicality in terms of such particularistic non-normative profiles entails strict identity. Note that properties like *being the author of Waverly* are repeatable, even though they identify a particular individual in the actual world. There are possible worlds in which someone other than Sir Walter Scott is the author of Waverly.

Let "Strong Repeatability" refer to the constraint that β includes only repeatable properties.[12] Given strong repeatability, we can say of any B in β that renders supervenience true for a normative property N that it is the case that $\Box_m \forall x \ (Bx \ \rightarrow \ \Diamond_m \ \exists y \ (By \ \wedge \ y \ \neq \ x))$.[13] Strong repeatability turns out to be a stronger constraint than what is needed for my point concerning expressivist explanations to go through. It suffices for the argument that the base excludes at least *some* non-repeatable properties, enough to ensure that there is some x and some y such that (i) $x \neq y$, (ii) x is B-identical to y.[14] We can refer to this weaker requirement as "Weak Repeatability." If the supervenience base is weakly repeatable, it excludes, e.g., the haecceitistic properties of the individuals that satisfy the above condition. Weak repeatability entails that the normative properties of at least some numerically distinct individuals supervene on their repeatable characteristics.

Discussions concerning supervenience rarely make the repeatability of the base properties explicit, but when we look at arguments for supervenience and the theoretical uses to which supervenience has been put in metaethics and ask which thesis metaethicists have had in mind, it looks to be repeatable supervenience in every case:

(i) The Inconceivability Argument

Consider the standard inconceivability arguments used to motivate supervenience. Given an individual with normative features, it is very hard to conceive of something identical in non-normative terms but different normatively. This imaginative exercise, which serves as the primary motivation for the conceptual truth of supervenience, standardly involves the inability to imagine that *two* non-normatively identical individuals differ normatively, *not* the inability to imagine two non-normatively identical individuals. The hypotheticals proponents of supervenience invite us to consider involve numerically distinct acts, persons, or states of affairs. For example, we're told Fred and Bob are exactly alike in all non-normative respects—they come into existence at the same time, have all the same experiences and dispositions, end up performing all the same actions—and we're invited to conclude that it is inconceivable for Fred to be a good person but not Bob. If such inconceivability intuitions can be treated as evidence for the conceptual truth of

supervenience, they represent evidence for the conceptual truth of (at least weakly) repeatable supervenience.[15]

In fact, such examples tend to show that the base excludes more than just non-repeatable characteristics. Some repeatable characteristics may be excluded as well, like spatiotemporal location, because it is implausible (though not entirely inconceivable) that Fred and Bob, who are supposed to be base-identical, could occupy the same space-time coordinates (they might be ghosts who have chosen to spend eternity occupying the same region of space). But the question of whether further limitations need to be imposed needn't concern us.

(ii) The Argument from the Ban on 'Mixed Worlds'

Blackburn (1971, 1993) relies on supervenience in a famous argument against cognitivism. He argues that the following two facts present an explanatory puzzle. Given a fully specified base property, B, for a normative property N, it is conceptually possible that something could be B but not N (NO CONCEPTUAL ENTAILMENT).[16] Yet it is inconceivable that there could be things in a single world that are B-identical but only one of them is N (WEAK SUPERVENIENCE / BAN ON MIXED WORLDS).[17] Blackburn suggests these two facts are puzzling if we accept a cognitivist realist account of normative concepts. What explains the fact that although there are conceptually possible worlds where B-things are N and conceptually possible worlds where B-things are not N, it is conceptually impossible for a single world to have some B-things that are N and some B-things that are not N? The result ostensibly violates plausible recombination principles for possible worlds and the conceptual constraint seems arbitrary.

We shall return to why Blackburn thinks this is a puzzle for cognitivists in particular in §4. For now, the key takeaway is that there would be no puzzle whatsoever if the supervenience base included all haecceities (violating weak repeatability). If B-identical things must be haecceitistically identical, then we have a straightforward explanation for the BAN ON MIXED WORLDS consistent with NO CONCEPTUAL ENTAILMENT: B-identical things are identical *tout court*, and one and the same thing in a world cannot be both N and not N (by the law of non-contradiction: $\forall x \forall \phi (\phi x \lor \neg \phi x)$). Since LNC only applies at a world, we don't violate NO CONCEPTUAL ENTAILMENT by explaining the BAN in this way. For the phenomenon to be at all puzzling, as Blackburn intended, it must be the case that B-identical things needn't be strictly identical—the base is such that it is possible for some genuinely distinct things to be B-identical.[18]

The assumption about repeatability should therefore be common ground between myself and my primary interlocutors: expressivists like Blackburn. In §3, we see why Gibbard (2003) is also best understood as having presupposed repeatability. It is especially important to the overall dialectic that expressivists have assumed the constraint.

(iii) The Argument from World-Bound Individuals

The final argument appeals to the role of supervenience as an important ex-planandum in metaethics. Suppose we read in tomorrow's newspaper that David Lewis (1968) was right: there is no such thing as strict or literal identity between individuals across possible worlds. When I consider myself in the nonactual possible world where I am a successful musician, I am not literally thinking about myself, but rather a very similar counterpart. Lewisian modal realists are not the only ones who find the notion of trans-world identity incoherent (Heller 1998, Sider 2001, Forbes 1982). But the reasons for thinking so are not our concern.[19] What is relevant is that no one ever understood the challenge of explaining normative supervenience in a way that turns on a proof that individuals are world bound. The conceptual truth of supervenience was supposed to function as a dialectical lever separating good metaethical theories from bad ones. For example, one of the main drawbacks of non-naturalist realism is the alleged difficulties faced by the non-naturalist in explaining supervenience (McPherson 2012). But if unrestricted (non-repeatable) supervenience were the relevant explanadum, then everyone, in-cluding the non-naturalist, should be relieved to discover that individuals are world bound, for we can explain unrestricted supervenience without assuming anything about the nature of the normative. We can simply appeal to the non-repeatability of the base properties plus the law of non-contradiction.[20] The reason it is a con-ceptual truth that it is metaphysically impossible for two base-identical things to be normatively distinct is that base-identical things are identical *tout court* and $\forall x \forall \phi(\phi x \vee \neg \phi x)$.[21]

The strength of this final argument lies in its showing not just that unrestricted supervenience is too weak, but that Weak Repeatability is precisely what we need to assume to arrive at the weakest non-trivial truth that could possibly serve the functions supervenience was intended to serve in metaethics. It is not like we can give up on the law of non-contradiction. And no one ever thought assumptions in modal metaphysics concerning identity across possible worlds would generate a simple explanation for supervenience that entirely bypasses consideration of the nature of the normative. Unrestricted supervenience is too easy to explain be-cause it allows all non-repeatable properties into the supervenience base, ensuring the literal identity of base-identical individuals.

So, the interesting thesis that metaethicists (and expressivists in particular) seem to have had in mind is that the normative supervenes on a repeatable non-normative base.[22] The above arguments don't quite establish full strength or 'Strong' Repeatability. But a good defense can be made for it on the ba-sis of Weak Repeatability, which the arguments do establish, combined with the assumption that whatever restriction on the base has been presupposed is likely to be a principled one.[23] The inference to Strong Repeatability is entirely optional, however, because my argument against the standard expressivist line on supervenience does not rely on Strong Repeatability. It relies only on Weak Repeatability.

3. The General Challenge Facing Expressivists

To deploy a normative concept, according to the expressivist, is not to predicate a property of a thing. On a rough characterization sufficient for present purposes, there are no normative properties on the expressivist account, not in any substantial sense at the very least (we can ignore quasi-realist wrinkles in what follows).[24] Though normative judgments, or utterances involving normative terms, share the surface grammar of ordinary descriptive judgments, they have a motivation-expressive function.

For ease of discussion, suppose the only normative concept on the scene is that of <being what ought to be done>, and this concept applies to act-circumstance pairs (ϕ, C) based on their features, where an act either is or is not what ought to be done in the relevant circumstance. According to the expressivist, the judgment that one ought to ϕ in C is wholly constituted by a motivation to ϕ in C. To assert that one ought to ϕ in C is to express the relevant motivation. For example, to judge that one ought to give a large portion of one's income to charity in circumstances of substantial inequality just is to be motivated to be appropriately generous with one's income in the relevant circumstances.

SUPERVENIENCE entails that it follows from the nature of the ought concept that it cannot apply differently to two act-circumstance pairs that are non-normatively identical (henceforth, base-identical or B-identical). Judging one ought to ϕ in C_1 but not in C_2 while recognizing the circumstances as B-identical would be akin to judging that Tim is a married bachelor. One can of course *assert* "one ought to ϕ in C_1 but not in C_2 and the two circumstances are B-identical." But the assertion would be unintelligible. The challenge is to explain why such judgments are unintelligible in terms of the special non-cognitivist semantics for normative terms (Gibbard 2003: 90–2; Compare: Klagge 1988, Blackburn 1984).[25]

The problem for expressivists is that such judgments appear intelligible if normative judgments are wholly constituted out of motivational states. Given repeatability, we know that C_1 and C_2 can be distinct circumstances—they can differ in some non-repeatable respect.[26] And one can plausibly and without conceptual confusion have motivational states directed at non-repeatable properties. For example, one can desire to ϕ in C_1 but not in C_2 *just because* they are distinct circumstances, a desire which constitutes a supervenience-flouting normative judgment. Likewise, one can care for an individual, and not merely as a bearer of repeatable properties. I may desire to treat Fred differently from Bob just because Fred is Fred and Bob isn't (rather than on the basis of Fred's repeatable qualities, like being a kinder person than Bob).[27] Such "identity-dependent" desires may be whimsical or morally problematic (Atiq 2016). But identity-dependent desires are not impossible. They can combine with one's beliefs in a situation to cause the agent to perform the desired action. So long as one happens to acquire true beliefs about which situation one is in (say, a situation involving Fred as opposed to Bob), one's desire to treat Fred differently from Bob will result in the desired outcome.

The nature of this challenge can be seen more vividly by considering that the expressivist has no trouble at all explaining unrestricted supervenience in terms of

constitutive constraints on motivation expression. If base-identical individuals are strictly identical, then a desire to act differently in base-identical situations amounts to a desire to act differently in a single situation. And it is true that one cannot desire to perform incompatible acts in one and the same situation. This is because no motivational state can cause one to ϕ and not ϕ in one and the same situation. It is possible, of course, to form *distinct* desires for a single circumstance that are jointly inconsistent (Sturgeon 2009).[28] After considering one's current situation at different times, one might form a desire to go out for dinner as well as a desire not to go out. But it is one thing to have two separate desires that happen to be incompatible and quite another to have a singular desire that aims at what one recognizes is an impossible end: ϕ-ing & not ϕ-ing in C_1. The latter is genuinely ruled out by *what it is* to desire.

So, while the expressivist can say that a judgment violating unrestricted supervenience is unintelligible because it cannot possibly involve the expression of a suitable desire, she cannot say the same thing about judgments violating repeatable base supervenience. And, as argued, the precise modal truth at stake in metaethics is repeatable supervenience.

4. Do Intentions (or Intention-like States) Provide Relief?

One response to this general worry is to observe that there are principled distinctions to be drawn between various natural kinds of motivations. Generic desires and mere preferences are essentially distinct from intentions, for example. Even if identity-dependent desires are possible, there may be constraints on intentions—constraints that derive from *what it is* to intend—that rule out an intention to distinguish situations based solely on their distinctness that are otherwise identical. So long as there are independent grounds for thinking expressivism is best developed in terms of such intention-like states rather than generic desires or preferences, a means of deflecting the challenge may be in the offing.

We'll return to the relevant constraints on intentions in a minute, but it is worth noting that something like this approach appears to be Gibbard's (2003). Gibbard thinks the relevant non-cognitive state is that of a *plan*. While planning is supposed to be an activity we all engage in and on which we have an intuitive grip, plans are, on Gibbard's construal, markedly flexible. We can plan not just for actual and contingent situations, but also for situations it is impossible for us to be in.[29] One plans to perform an action ϕ in a circumstance C, so that ϕ-ing in C can be described as the plan's aim.[30] There can be grander plans, ones that specify actions for multiple circumstances, and even a universal plan which specifies an action for every possible situation: a plan to ϕ_1 in C_1, ϕ_2 in C_2, ϕ_3 in C_3, and so on.

As one would expect, Gibbard thinks supervenience-flouting normative claims are unintelligible because a certain kind of plan is impossible: a plan to ϕ in C_1 and not ϕ in C_2 while recognizing C_1 and C_2 as B-identical (2003: p. 93–8). The plan isn't just one that our community of planners generally avoids or dislikes. It is impossible, and this is supposed to be knowable *a priori* based on introspective awareness of one's own mental states. If such a plan were possible, given the metasemantics, one

could intelligibly judge that one ought to ϕ in one circumstance but not the other while recognizing the circumstances as B-identical. Since the plan is impossible, the corresponding judgment is unintelligible.

Gibbard does not consider the repeatability of the supervenience base (explicitly, at any rate). On Gibbard's construal, the supervenience base is the class of scientifically observable properties, or as he sometimes puts it: "prosaically factual properties" (p. 32, 98–99). The category is meant to include not just properties currently recognized by the natural sciences, but "spooky properties like the properties of ghosts and gods, so long as these can be recognized or observed and thus could figure in an empirical science."[31] Presumably, Gibbard intended the class to exclude haecceitistic properties of individuals and situations, given that such properties are unlikely to be scientifically observable. Either way, what matters is that the supervenience base *should* be defined in terms of repeatability if the aim is to explain the precise modal claim at stake.

So, the question for plan-based expressivists put more sharply is: why is a plan to act differently in two observably identical situations based solely on their distinctness impossible? Gibbard notes:

> a plan can distinguish between situations only in terms of their prosaically factual properties, and it can distinguish between acts only in terms of the prosaically factual properties of those acts. If two acts in two possible situations differ in no prosaically factual way, a plan cannot distinguish them, permitting one and ruling out the other. Either it will permit both or it will rule both out (2003: 91–2).

The argument for the claim goes by very quickly, but the basic idea is that outcomes alike in prosaically factual respects are quite simply indistinguishable from the planning perspective.[32] Planning to ϕ in C_1 but not ϕ in C_2 when there is no way of telling C_1 apart from C_2 based on observable characteristics would be just like planning to ϕ and not ϕ in what one recognizes to be one and the same situation. And a putative plan to ϕ and not ϕ in a single situation would be no plan at all.

If unrestricted supervenience were the relevant expanandum, Gibbard's claims would be quite compelling and we wouldn't have to worry about distinguishing plans from motivational states generally: a plan to perform incompatible acts in base-identical situations would be equivalent to a plan to perform incompatible acts in one and the same situation. And no motivational state can cause one to perform incompatible acts in one and the same situation. However, once we see that, given repeatability, base-identical situations needn't be strictly identical, the parallel Gibbard takes for granted becomes highly questionable.[33] A desire to perform incompatible acts in two genuinely distinct situations just because they are distinct is quite possible. Since it isn't generally the case that our motivations are constrained to be insensitive to unobservable features like situational haecceities, why should we think plans are so constrained?[34] (For now, we can set aside the possibility of understanding the plan-based expressivist as stipulating an artificial notion of a plan with the constraint built in. We'll return to why brute stipulation is not a viable strategy in §4.)

The plan-based expressivist must motivate the alleged constraint on planning, and one way to do so is to construe plans as closer to intentions than ordinary desires or preferences. There are independent reasons for thinking intentions (and intention-like states) constitutively involve not just a favoring of some outcome but an expectation, construed as reasonable from the point of view of the agent, that she will realize the outcome aimed at. Unlike generic desires, intentions involve taking for granted that one will do as one intends. At the very least, one cannot intend to realize an outcome one believes to be exceedingly unlikely. It is this reasonable expectation of a likelihood of desire-realization that is missing in the case of a desire to treat two observably indistinguishable situations differently *just because* they are distinct situations. The only possible worlds where this identity-dependent desire is satisfied involve luck—worlds where one has randomly acquired true beliefs about which of the two observably indistinguishable situations one is in.

But there is a fairly decisive objection to this strategy. We must keep in mind that whatever motivational state underwrites normative thought and talk, it must be flexible enough to capture the full range of meaningful normative claims. And there are clearly meaningful normative judgments concerning what one ought to do which cannot involve the expression of a simple intention to do it or, for that matter, the expression of a motivational state which constitutively involves subjectively rational expectations of desire realization. For example, expressivists are (and indeed should be) committed to vindicating as meaningful those of our ordinary normative judgments which involve so-called robustly objective uses of ought. Such usage involves judging that a person ought to do something even when the person does not know and could not know that it is the thing she ought to do. The person might be constitutionally disposed to doing the wrong thing.[35] Consider:

> <Paranoid schizophrenics ought to get themselves some therapy>

According to the expressivist, in asserting this fact I express a motivation to get therapy in the circumstance of being a paranoid schizophrenic. But if I were a paranoid schizophrenic, I might be insensitive to facts about my own paranoia and have no inclination to get therapy. If my psychological propensity to not ϕ in a situation because of my inability to reliably identify the situation based on my evidence makes my intending to ϕ in that situation impossible, the expression of such an intention cannot be a condition on meaningful normative assertion. Otherwise, we would have to give an error-theoretic diagnosis of what seems to be a clearly intelligible and quite likely true normative claim: that paranoid schizophrenics should get therapy.

Here's another example:

> <Joe ought to turn the light switch on and off repeatedly if doing so will, unbeknownst to him, cure cancer>

While it is controversial as a matter of first-order moral theory whether such objective ought claims are true, they are at least meaningful. Although one can desire to turn the light switch on and off in a circumstance where doing so will unknowably

cure cancer, one cannot intend to do so, given that one cannot realistically expect to realize the desired outcome.

Accordingly, expressivists must either accept that the non-cognitive attitude involved in normative judgment is closer to ordinary desires or preferences than full-blown intentions in that the association between the relevant desire-like state and expectations of desire-realization is contingent rather than constitutive; or else find some means of finessing objective ought claims. One approach involves appealing to the intentions of an appropriately idealized agent: say, one who is perfectly omniscient with respect to Joe's situation.[36] Call the agent "Hera." To judge that Joe ought to turn the light switch in his situation would be to express how one would intend were one Hera in Joe's situation. Because Hera is ideal, she has knowledge that the arbitrary act cures cancer and so will intend to perform the act.

One problem with this view is that it is entirely nonobvious what a perfectly omniscient being would intend for Joe's situation. To get around this problem, the idealization can be constrained to take into account only those features of the situation that the person making the normative judgment deems relevant (in the above case it would be the fact that the light switch is unknowably linked to curing cancer). There may be other problems with the view, helpfully explored in Daskal (2009)—e.g., that it makes it hard to make sense of disagreement between those who treat different features of situations as normatively relevant. But let us suppose it can be made to work.

The key point, given our purposes, is that if such moves are permitted, it becomes very hard to see why we can't similarly make sense of supervenience-flouting normative claims. If I can express Hera's intentions—an agent who can know the cancer curing facts in the light-switch case, I should be able to express the intentions of a similarly ideal agent, who knows all the haecceitistic facts about a situation. There is no conceptual barrier to a haecceity detecting ideal agent, one who can distinguish Fred from Bob based solely on their respective haecceities and even if Fred and Bob are observably indistinguishable. And so, to be consistent, our imagined expressivist who allows idealization should say that in judging that Fred and Bob normatively differ solely because Fred is Fred and Bob isn't, I express the intentions I take the ideal (all-knowing) agent to have in situations involving Fred and Bob (an intention to treat Fred and Bob differently based solely on their haecceities).

Couldn't the expressivist simply stipulate that moralizing involves idealizing with respect to some unobservable properties but not with respect to haecceities? No. And the reasons why brute stipulation won't work are discussed shortly in §4. For now, the key point is that if we must treat as meaningful normative claims which turn on features of a situation that the adjudicator knows she wouldn't be in a position to detect were she herself in the situation (by invoking a suitably ideal agent who would be in a position to detect the relevant features), it would be unprincipled not to treat as equally meaningful normative claims that turn entirely on the haecceitistic features of a situation.

Here's a different way of putting the point. Intentions seemed a promising natural kind of mental state to invoke in an explanation of supervenience because of

constitutive constraints that prevent us ordinary agents from forming intentions to distinguish observably indistinguishable situations based solely on their haecceitistic features. But those very same constraints prevent us from intending to φ in a situation C based on features of C we cannot detect. This makes it hard to make sense of clearly meaningful objective ought claims. If, in order to get around this new problem, we must invoke the intentions of an ideal agent who can detect features that we ordinarily cannot, then such an agent should be able to detect haecceities, which prevents us from ruling out as meaningful supervenience-flouting normative claims. And we're back to square one.

5. Can the Conceptual Truth be Saved by Stipulation?

Let us recap. The expressivist tells us that meaningful normative assertion involves motivation expression. We observed that there aren't any general constraints on motivation that preclude identity-dependent motivation. The expressivist responds that there are distinct natural kinds of motivations, some of which do appear to preclude pure identity-dependence on account of their nature, because they essentially involve a subjectively rational expectation that the desired outcome will be realized: intentions, for example. But, we observed, it cannot be a condition on meaningful normative assertion that one simply express such intention-like states, because it would rule out the meaningfulness of judging that one ought to do the right thing in circumstances involving not knowing what the right thing to do is. In short, we have been given no reason by expressivists to think supervenience-flouting normative judgments are unintelligible.

It turns out to be very difficult to find a natural kind of motivational state that can underwrite the full range of meaningful normative claims, while also being constitutively insensitive to the identities of things. No expressivist to my knowledge has succeeded in identifying such a state. In light of this, it might seem tempting to make do with an artificial category of motivations. We can define up a category of desires* which includes only identity-independent desires. We can then stipulate that normative concept use necessarily involves the expression of desires*, so repeatable supervenience-flouting judgments are unintelligible. Analogously, one might define up the concept of a max-desire, which is a lot like the concept of a desire, except that one max-desires to ϕ only if ϕ-ing involves maximizing happiness. It is impossible to max-desire not to maximize happiness because of what it is to max-desire.

This stipulative approach seems attractive initially, but there are several reasons for rejecting it. For one, it would render normative concepts and whether one succeeds at deploying them arbitrary. While I can't desire* to act differently in situations solely on account of their distinctness, I can do something which looks a lot like desiring*—namely, I can desire to act differently. It is hardly plausible that whether I succeed at meaningfully deploying normative concepts can hang on so tenuous a thread.

Second, desires seem to be the more natural or joint-carving category, whereas desires* are a gerrymandered class. There are good reasons to think that the meanings of our terms should respect objective joints in nature (Lewis 1983, 1984), and,

accordingly, good reason to think that if normative judgment involves motivation expression, the relevant motivations represent a natural kind. Another way to put it is that there is something unique about normative judgments in virtue of which they are *normative*, and desires (rather than desires*) represent the more natural joint at which to carve out any entirely distinct category of judgment.

Third, normative concepts aren't amenable to stipulative definition. The expressivist cannot simply stipulate that our ordinary concept of the good involves desiring* and not desiring. The claim must be motivated. And, crucially, the expressivist cannot appeal to our commitment to SUPERVENIENCE as proof that moralizing involves desiring* because an explanation of our commitment to SUPERVENIENCE was supposed to emerge from reflecting on the nature of normative concepts. That is, we were supposed to have a more explanatorily basic grasp on the nature of the non-cognitive attitude expressed when we use normative terms like 'good' and 'ought' than on our commitment to SUPERVENIENCE. And if there is no other justification expressivists can give for limiting meaningful normative judgments to those that involve the expression of a desire*, the restriction will seem *ad hoc* and designed to save the view from objection.

It is worth emphasizing this last point as a summary of the dialectic so far. We must keep in mind the order of explanation: our verdict on supervenience is supposed to be explained by our understanding of normative concepts and not the other way around. Since on the expressivist picture an understanding of normative concepts turns out to be an understanding, perhaps implicit, of the non-cognitive mental state we express when we deploy normative concepts, it had better be the case that we know we are desirers* and not desirers (when we deploy normative concepts) prior to our acceptance of SUPERVENIENCE. But we don't know this. Insight into what we're up to when we moralize does not provide *any* independent support for supervenience understood as a conceptual truth.

Furthermore, if normative judgments involve the gerrymandered category of desires* as opposed to the more natural category, presumably it is in virtue of facts about our usage of normative terms. Why, one wonders, would we have a linguistic convention which involves rejecting as meaningless attempted uses of normative language to express (identity-dependent) desires? One option is to insist on its being a brute fact. Perhaps there is no deeper explanation for why normative talk works this way. As far as brute facts go, this one needn't be all that mysterious. But if the conceptual truth of supervenience is to be explained by appeal to a perfectly arbitrary stipulation made by moralizers to make supervenience a standard for conceptual competence, then there is nothing distinctively expressivist about the explanation anymore. *Any* theorist can invoke such an arbitrary choice (Cf. Jackson 1998, p. 125; see discussion in Dreier 1992, p. 21), and one of the advertised virtues of expressivism was precisely that it can avoid such arbitrariness.

It is indeed telling that expressivists have long claimed the ability to provide a deeper and distinctively expressivist explanation of the conceptual truth of SUPERVENIENCE.[37] Expressivist explanatory ambitions have presumably been guided by the attractive assumption that the line between meaningful and meaningless normative claims does not reflect an arbitrary choice amongst moralizers. Blackburn

(1971, 1993), for example, insists that the conceptual truth of supervenience supports expressivism because he finds the appeal to an arbitrary choice in the case of normative concepts unsatisfying. He admits that any cognitivist can invoke such a choice but finds it puzzling that moralizers would have adopted a supervenience constraint for no reason whatsoever, rather than some other conceptual constraint or no constraint at all, given that (for broadly Humean reasons) there is no conceptual entailment from non-normative facts to normative ones (see discussion of the Argument from the 'Ban on Mixed Worlds' in §1).

Blackburn may be right about the unsatisfying nature of an appeal to arbitrary stipulation, but he is wrong to think expressivists can avoid this problem, at least if expressivists are forced to explain the conceptual truth of supervenience by stipulating that the desires involved in moralizing are supervenience-constrained. And, for what it is worth, I think Blackburn is right about the problems with brute conceptual stipulation. In any event, while it might be true that our normative linguistic conventions are more plausibly taken as arbitrary by an expressivist rather than, say, a non-naturalist realist about the normative,[38] the stipulative move remains less than satisfying for anyone who took it to be an important pre-theoretical datum that the truth demands explanation.

Instead of treating the convention of making supervenience a standard for meaningful normative assertion as arbitrary, the expressivist might try to provide some rationalizing explanation for it. Perhaps a richer characterization of the essential functions of normative discourse will reveal reasons why we don't count the expression of identity-dependent desires as meaningful normative assertions. Consider Blackburn's (1986) own account of the functions of normative language. According to Blackburn, we use normative language to express our stable, higher-order desires with the aim of coordinating with one another. We engage in normative discussion, and the associated expression of our desires, to invite others to share our motivations so we can harmonize our ends. The coordination aim might very well explain our disposition to negatively evaluate those who express supervenience-flouting motivational states. It would predict that we would be inclined to use normative language to express our disagreement with those who reject supervenience insofar as our own motivations are supervenience-respecting. But it is hard to see why the coordination aim of normative language should entail that normative claims underwritten by supervenience-flouting non-cognitive attitudes would be rejected as *meaningless*.[39] Normative language can facilitate coordination amongst agents even if it involves the expression of stable desires (the more natural category) rather than stable desires that are necessarily identity-independent.[40]

This does not amount to a decisive showing that there could not be a functional explanation for why normative language is essentially tied to the expression of supervenience-respecting motivations. But the burden falls on expressivists who wish to pursue this strategy to come up with a more plausible explanation. In the absence of one, our default position should not be that there is an undiscovered functional explanation in the offing.

6. The Normative Alternative

Fortunately, expressivists are not all out of options at this point. There is an alternative approach to supervenience that warrants serious consideration in light of the weaknesses of the traditional approach. It involves treating repeatable supervenience as an ordinary normative truth, and explaining it in the way that expressivists generally explain normative truths.

To assert supervenience is to *express a desire* to not treat situations differently based solely on their haecceitistic features. To judge that the mere distinctness of things could not possibly make a normative difference is to make a paradigmatic first-order normative claim. It needn't be taken as fundamental. The expressivist can normatively explain why we might reject identity-dependent desires. Such desires lead to the frustration of one's aims, for we are unable to reliably track the haecceities of individuals. Moreover, the identities of agents or situations do not seem to matter from the impartial point of view (Atiq 2016). In other words, there is a distinctively moral explanation in terms of the equality of persons and the good of treating others impartially for having identity-independent desires. This approach allows the expressivist to give a deeper explanation for the truth of supervenience, albeit a normative one, consistent with our guiding assumption that the truth calls out for explanation and shouldn't be taken as brute. The normative approach has the further advantage of avoiding arbitrary metasemantic stipulations.

There is a minority position in metaethics according to which SUPERVENIENCE should be treated as a non-conceptual normative truth (Fine 2002; Kramer 2009; Rosen forthcoming). Expressivists who opt for this strategy may be able to draw additional support from the general considerations these theorists offer for their normative characterization of SUPERVENIENCE.[41]

But whether or not the expressivist is in a position to rely on these general considerations, expressivists have sufficient theoretical resources of their own to explain why supervenience might have seemed to be a conceptual truth. For one, the normative truth of supervenience (like conceptual truths generally) has a distinctively *a priori* character given that all it takes to recognize it, given expressivism, is to have the relevant impartial or identity-independent motivations. Recognition of the normative truth does not turn on subtle empirical observations.[42] Second, the expressivist can appeal to our difficulty imagining alien motivational orientations to explain why we tend to find denials of supervenience incoherent. Preferences that treat individuals differently *just because* they are distinct individuals are highly unusual (contrast a preference for a particular twin based on some repeatable characteristic like *being the kinder twin*).[43] And it is difficult to imagine what it would be like to have unusual and morally arbitrary preferences. Compare the difficulty in making sense of someone who thinks genocide is permissible. The truth that genocide is impermissible has the same conceptual seeming character, and, arguably, for similar reasons: it is impossibly difficult to imagine being motivated to commit genocide. Our imaginative failures that come in the way of comprehending a supervenience-flouting normative judgment are just that: failures of imagination (that reflect well on us morally, one might add).

For all the advantages of the normative approach, there is a dialectical cost that expressivists unavoidably incur by denying supervenience the status of a conceptual truth, but one that is not worth losing sleep over.[44] Expressivists have long regarded it a central virtue of their account that, unlike competing views on the nature of the normative, expressivism is uniquely compatible with and explains the conceptual truth of normative supervenience. For example, Blackburn (1993: 137) writes that it is a principle virtue of expressivism that it fully explains why "it seems to be a conceptual matter that moral claims supervene . . . . Anyone failing to realize this or to obey the constraint would indeed lack something constitutive of competence in the moral practice." Having to take repeatable supervenience to be an ordinary normative truth undercuts the expressivist's right to claim an explanatory advantage over other views that has long been taken for granted. Fortunately, there are other, better reasons to be an expressivist than any explanatory advantages concerning supervenience (Atiq MS); and no reason not to be an expressivist just because one has to deny that supervenience is a conceptual truth, because the normative characterization fits our intuitions about supervenience just as well.

Notes

[1] There is a minority position which rejects the conceptual status of the truth of supervenience (Fine 2002; Kramer 2009; Rosen forthcoming). I discuss its significance in §5.

[2] I distinguish weaker and stronger versions of this constraint. If the supervenience base is weakly repeatable, then some (but not all) non-repeatable properties are excluded, enough to ensure that the normative properties of at least some numerically distinct individuals supervene on their repeatable features.

[3] Doubts about expressivist accounts of supervenience have been raised before (Zangwill 1997, 110–11; Sturgeon 2009, 84–87; Dreier 2015). But the vulnerability which concerns me and that I distinguish in §§2 and 3 has been overlooked.

[4] It is easy to multiply examples of theorists across the metaethical spectrum who endorse the thesis. See e.g. Smith (2004: 208), Blackburn (1984), and Enoch (2011).

[5] There are weaker versions of the supervenience thesis, concerning the co-variance of normative and non-normative properties within possible worlds (Kim 1984). But we will mainly be occupied with the strong formulation in what follows.

[6] This formulation assumes non-normative properties are closed under infinitary Boolean operations. Since reductive naturalists deny this, it's not an entirely neutral formulation. One way of avoiding this problem is by talking about sets of normative properties (Bader 2016). But that still does not get the reductive naturalist on board who identifies a normative property with a simple non-normative property (like *being pleasing*), supposing that a property cannot be both normative and non-normative. For a helpful discussion of how to accommodate such views, see McPherson (2012). My argument does not turn on these complications.

[7] One reason it might have seemed tempting to leave the class unrestricted is to get the moral particularist on board. Moral particularists deny that there are exceptionless general principles linking the normative to the non-normative (Dancy 2004). It is true that my argument for constraining the supervenience base rejects an extreme kind of moral particularism, according to which even the haecceities of situations and possible worlds could matter from the normative perspective. But if I'm right, we shouldn't have been concerned with accommodating extreme particularism. The supervenience claim I defend leaves room for more modest versions of particularism on which there are no snappy general principles in ethics, but we know that some kinds of non-normative properties are ethically irrelevant.

[8] This particular problem won't arise if the property of *being the actual world* is normative, supposing, e.g., that its real definition includes reference to normative truths at the world. But world-indexed

properties are not obviously normative. I consider three additional arguments below for why the base can't be left unrestricted which do not rely on the assumption that the base includes world-indexed properties.

[9] As Blackburn points out, too much precision in specifying the relevant non-normative properties would render the claim synthetic: possibly true, but not true in virtue of the concepts involved. Consider the debate between deontologists and utilitarians about whether an act's rightness covaries with the non-normative property of being happiness maximizing. Their disagreement does not appear to be a conceptual one. Nevertheless, Blackburn does assume some minimal conceptually-derived constraint on the supervenience base, as I discuss below.

[10] The supervenience base for the normative was traditionally characterized in terms of natural properties, where a property counts as natural if it is invoked by the natural sciences (microphysical and causal/functional properties, for instance) or if it would be invoked by the sciences sufficiently idealized (Jackson 1998, Smith 2004; Cf. Griffin 1996). It is sometimes suggested that all the familiar ways of defining the supervenience base render the thesis vulnerable to counterexample (Sturgeon 2009). I take these skeptical arguments to show only that stating the claim precisely is hard, not that supervenience is false.

[11] Repeatable supervenience is entailed by Hare's (1952, 129) claim that moral judgments are necessarily general and universalizable:

> As we shall see, all value-judgements are covertly universal in character, which is the same as to say that they refer to, and express acceptance of, a standard which has an application to other similar instances. If I censure someone for having done something, I envisage the possibility of him, or someone else, or myself, having to make a similar choice again; otherwise there would be no point in censuring him . . . . When we commend an object, our judgement is not solely about that particular object, but is inescapably about objects like it.

Hare thought that for a value judgment to be truly universal, it cannot include proper names (p. 176–177). Repeatability allows that morally relevant non-normative properties might involve essential reference to particulars using proper names. E.g., the property of *being a kicking of Fred* is repeatable (there are many acts that are kickings of Fred), though a principle which treated the property as ethically relevant would not be suitably universal on Hare's view. What repeatability rules out is the inclusion of *being Fred* in the supervenience base.

[12] What about prudential obligations to oneself? If I'm obligated to care about myself simply because of my identity, is this not a counterexample to 'strong' repeatable supervenience? The obligation to care about oneself is not necessarily identity-dependent. The relevant non-normative fact which grounds my obligation to myself is that I am a person with interests (one whose interests I am intimately familiar with), but *being a person with interests* is repeatable. Even if prudential normativity is sensitive to identity-based facts, my point about expressivism goes through so long as there are some normative domains where haecceitistic properties do not matter, or some haecceitistic properties that are normatively irrelevant.

[13] There are two ways of implementing this constraint. Either β, the set of base properties, is constrained so that $(\forall B \in \beta)\ \Box \forall x\ (Bx \rightarrow \Diamond\ \exists a\ (Ba \wedge a \neq x))$. Alternatively, we can include the constraint in the statement of supervenience: $\Box_C(\forall N\ \text{in}\ \alpha)(\forall x)[Nx \rightarrow (\exists B\ \text{in}\ \beta)(Bx\ \&\ \Box_M(\forall y)(By \rightarrow Ny)\ \&\ \Diamond_M\ \exists a\ (Ba \wedge a \neq x))]$.

[14] If the base is weakly repeatable, we know that in addition to the general SUPERVENIENCE claim being true, $\exists a\ (\forall N \in \alpha)[Na \rightarrow (\exists B \in \beta)(Ba\ \&\ \Box_M(\forall y)(By \rightarrow Ny)\ \&\ \Diamond \exists b\ (Bb \wedge a \neq b)]$.

[15] Here's a related argument for thinking the type of inconceivability which motivates supervenience involves consideration of two genuinely distinct normative property bearers. Suppose that in trying to imagine individuals that are non-normatively identical we can only imagine a single self-identical individual. This raises the worry that what we really intuit when it seems to us we cannot normatively distinguish individuals after fixing their non-normative features is not the truth of supervenience but the logical tautology that a single individual cannot both be N and ¬N. The imaginative exercise is more probative of the truth of supervenience if it involves imagining two genuinely distinct acts that are non-normatively identical and discovering that these distinct acts must be normatively identical.

[16] Even if N supervenes on B, Bx does not conceptually entail Nx, for all x, N, & B. Blackburn thinks this for familiar Humean reasons (the is-ought gap) and the open question argument. Note

that No CONCEPTUAL ENTAILMENT is compatible with SUPERVENIENCE, because the nested modality is interpreted as metaphysical necessity not conceptual necessity. SUPERVENIENCE entails only that the conceptually necessary truth is that there is some B (we don't know which) such that Bx entails with metaphysical necessity Nx.

[17] Blackburn describes it as a ban on mixed worlds in *Spreading the Word* (1984, p. 184). He thinks this is true in the case of *actual* non-normative qualitative duplicates: two states of affairs in the actual world, one located in the northern hemisphere and the other in the southern hemisphere, which are otherwise non-normatively alike must be normatively identical. And we can know this, according to Blackburn, without knowing which non-normative properties the two states exemplify.

[18] As it turns out, when Blackburn revisits his discussion of the ban, he's quite explicit that *some* restriction on the base must be assumed. In *Supervenience Revisited* (p. 133), he notes:

> First, we do not want the supervenience thesis to be made vacuously true through its being impossible that any two distinct things should be G* [identical in base respects]—it then following that if one G* thing is F they all are. And, the threat here is quite real. To get around this I am going to assume a limitation thesis. This will say that whenever a property F supervenes upon some basis, there is necessarily a boundary to the kind of G properties that it can depend upon.

Blackburn does not specify the nature of this assumed boundary. I've tried to show the boundary *must* be such that for some x and some y, their normative properties supervene on their repeatable features. And the reason I've focused on this limitation is that it is all I need to challenge the standard expressivist line on supervenience.

[19] One might think this because of a commitment to modal realism: one and the same concrete individual cannot be part of two entirely distinct concrete worlds. Alternatively, one might be a type of ersatzer who thinks only actual individuals have individual haecceities, but denizens of merely possible worlds (merely possible individuals) have no such individual essences.

[20] We need to assume individuals are world-bound—i.e. no strict or literal identity between individuals in different possible worlds—because LNC only applies to individuals at a world. Of course, not everyone rejects transworld identity, but the point of the argument still stands: supervenience is too easy to explain if B-identity entails strict identity. It can be explained by appeal to (1) B-identity = strict identity; (2) LNC: $\forall x \forall \phi (\phi x \lor \neg \phi x)$; (3) Worldbound individuals. Thanks to Harvey Lederman for pushing me to make this explicit.

[21] An anonymous reviewer objects that because the explanation relies on the assumption of world-bound individuals it wouldn't necessarily explain why supervenience is a *conceptual* truth. The objection doesn't work for several reasons. First, we can invoke WBI in an explanation of the conceptual truth of supervenience because WBI may well be a conceptual truth. At the very least, a proponent of WBI might reasonably think this. Arguments for the unintelligibility of transworld identity typically invoke facts about the concept of identity. A familiar line of argument stresses that the concept of identity is the concept of a relation grounded in causal continuity and the like, and it is very obscure what it would mean for such a relation to obtain across possible worlds (see e.g., discussion in Hazen 2012, pp. 46–7).

Second, even if WBI is not a conceptual truth, this does not mean it can't play a role in an explanation of SS as a conceptual truth. Conceptual explanation often involves assuming background *a priori* principles that don't appear to be analytic. E.g., Williamson (2006) has a nice example of explaining <every vixen is a female fox> to someone who has considered philosophical reasons for embracing a logic on which universal generalizations are existentially committing and happens, also, to have a false empirical belief that there are no vixens. This individual reasonably wouldn't accept an explanation of the ordinary conceptual truth based solely on the intersubstitutability of "vixen" and "female fox." Analogously to logical principles, WBI may be an *a priori* non-conceptual truth that is part of the background framework for a conceptualist explanation of SS. Note that despite assuming WBI, we still have to invoke facts about the concepts in question—that the supervenience base for the normative is not repeatable (ex hypothesi)—to explain unrestricted SS.

The reviewer points out that it strains credulity to think those who endorse supervenience as a conceptual truth have had WBI in mind as their reason for doing so, given the nonobviousness of WBI. Perhaps this means an explanation that relies on WBI must be mistaken. But the objection proves too much. Consider an expressivist who attempts to explain supervenience by appeal to the

motivation-expressive metasemantics of our ought concept and constitutive facts about the nature of our motivations. It strains credulity to think those who endorse supervenience as a conceptual truth have had these expressivist insights in mind. After all, many proponents of SUPERVENIENCE emphatically deny expressivism. If the objection works, *no one* can explain SUPERVENIENCE by appeal to a controversial nonobvious theory of normative concepts. We avoid this result by admitting that individuals may intuit the truth of a claim without seeing clearly why it is true.

[22] Repeatable supervenience is clearly true in many areas. If two things are exactly alike in all repeatable underlying respects, either both are chairs or neither is. Why is that? It is somehow a conceptual truth that what makes a chair a chair is a matter of its repeatable features, including its shape, its purpose, and so on. What seems true in the case of the chair concept is also true in the case of the concept of goodness. Thanks to Gideon Rosen for the comparison.

[23] Here's how we might get to the stronger claim. My arguments show that the supervenience thesis presupposed excludes at least *some* non-repeatable characteristics from the base. It is hard to come up with a non-arbitrary reason for excluding some but not all non-repeatable properties. Moreover, it is intuitive that haecceitistic and other non-repeatable features are normatively insignificant (it's hard to come up with a counterexample). So, presumably the widely assumed constraint on the supervenience base is: n*o* non-repeatable properties. The class won't be fully closed under Boolean operations, but this technicality needn't concern us. An anonymous reviewer suggests the non-repeatable property of *being the unique necessary being* may be a counterexample to Strong Repeatability. I find it hard to see why a being might be praiseworthy or valuable simply in virtue of being the unique necessary being. Contrast features like benevolence or omniscience which might justify praise. While it is true that uniqueness in general seems possibly valuable (cf. Helen's unique way of being beautiful), the exemplification of some way of being unique amounts to the exemplification of a property that is repeatable.

[24] Quasi-realist expressivists are happy to admit talk of normative facts and properties, but with a deflated account of such talk. To judge it is a fact that one ought to help the poor is equivalent to judging that one ought to help the poor. Both judgments receive the motivation-expressive treatment. My argument does not turn on whether we assume a quasi-realistic version of expressivism. Whether quasi-realists confront a unique explanatory challenge is helpfully explored in Dreier (2015).

[25] For a careful derivation of the unintelligible judgment that stands in need of explanation, see Klagge (1988).

[26] Weak repeatability is all we need to assume here. There is some C_1 and some C_2 such that C_1 and C_2 are normatively identical and their normative properties supervene on their repeatable characteristics.

[27] A metaphysician who believes in haecceities could have desires or intentions directed towards the alleged haecceities of certain individuals.

[28] Sturgeon (2009) uses such cases to raise a challenge for Gibbard's plan-based expressivism (and expressivism more generally). Sturgeon suggests that a commitment to SUPERVENIENCE involves accepting a rational norm of revising pairs of normative judgments jointly incompatible with SUPERVENIENCE. Sturgeon (2009) demands an explanation for why discovering inconsistencies in one's plans (or one's motivations generally) should be universally recognized as "a cause for alarm, or reform . . . rather than just for amusement?" My purpose in raising Sturgeon's objection is not to criticize or defend it, but to note that the case he describes (because it involves two distinct motivations) does not amount to a direct counterexample to the conceptual truth of supervenience given expressivism—that is, a case of a supervenience-flouting normative judgment that the expressivist's metasemantics validate as intelligible.

The challenge confronting expressivists is more direct than critics have realized. It does not require assuming that a commitment to being consistent falls out of our acceptance of supervenience as a conceptual truth. The basic challenge for expressivists consists in the fact that there are intelligible, naturally occurring pro-attitudes that distinguish non-repeatable characteristics like identity. And these pro-attitudes constitute intelligible normative judgments that violate supervenience if expressivism is true.

[29] It is crucial for plans to be flexible in this way if the plan-based expressivist is to be able to capture the full range of normative judgments. I return to the intended flexibility of plans in the discussion below.

[30] Technically, Gibbard uses sets of what he calls 'fact-plan worlds' to interpret states of mind expressed, where a fact-plan world can be thought of as a pair (f, p), consisting of a state of the world f and a plan for that state. A fact-plan world (f, p) is compatible with the state of mind expressed when

you judge that you ought to phi in a given situation if the 'fact' component is consistent with what you believe the world is like, and the second element—the plan—has you phi-ing. The state of mind is then characterized in terms of the set of all fact-plan worlds compatible with one's view of the world and one's motivations regarding what to do. These complexities can be ignored in what follows.

[31] The relevant notion of empirical recognizability is not altogether clear. In referring to the properties as "prosaically factual," it seems Gibbard intends the notion to be quite strong. For example, it might involve phenomenological recognizability, where if a property is recognizable it is directly available in phenomenal experience or else impinges in some way on our sensory surfaces. More plausibly, the category includes just those properties whose existence can be taken for granted under the norms of ordinary scientific practice.

[32] Dreier (2015) interprets Gibbard similarly: "It is of the nature of plans that the contingencies for which they are plans are features of the situation that are at least in principle recognizable."

[33] Gibbard's intriguing claims about the nature of planning have been curiously neglected. Hawthorne (2002: 173–4) expresses some skepticism, though only in passing.

[34] For all Gibbard tells us about plans (namely, that plans are action-guiding motivational states individuated using action-circumstance pairs), they may be just like desires or preferences.

[35] See Daskal's (2009) discussion of Gibbard's view concerning what Gibbard should say about 'Mortis,' the ideally coherent psychopath.

[36] I credit an anonymous reviewer for suggesting this approach.

[37] Blackburn (1970: 120; 1984: 64): "[Supervenience] is *explained* by the anti-realist nature of moralizing." See Gibbard (2003: 96) and discussion of Gibbard's view in §3.

[38] If the non-naturalist takes it to be a brute fact about our linguistic practice that we do not recognize normative judgments inconsistent with supervenience as meaningful, she confronts the question: how is it that an arbitrarily determined feature of our practice mirrors a robustly metaphysical fact in normative reality—namely, the supervenience of robustly real normative properties on repeatable non-normative properties?

[39] So, when Blackburn (1993) writes:

> it seems to be a conceptual matter that moral claims supervene on natural ones. Anyone failing to realize this or to obey the constraint would indeed lack something constitutive of competence in the moral practice. And there is good reason for this: it would betray the whole purpose for which we moralize, which is to choose, rank, approve, for bid, things on the basis of their natural properties. (p. 66)

he presupposes the fact that needs to be explained—namely, that the purpose behind using normative terms (or moralizing) is to choose, rank, approve and so on *specifically on the basis of the natural properties of things* (construed as a repeatable supervenience base). The puzzle is precisely to explain (or provide some independent support for) the claim that when we moralize we aim to coordinate on the basis of repeatable non-normative properties alone. (As argued in §1, Blackburn should be construed as assuming repeatability, based on his arguments concerning supervenience and the ban on mixed worlds.)

[40] Two metaphysicians who believe in haecceities may wish to coordinate over their shared fetish for the identities of things.

[41] Fine and Rosen deny that supervenience is a conceptual truth because they don't even think it's true (if the nested modality is interpreted as metaphysical necessity). They don't think it's true because of their commitments concerning metaphysical necessity: they think such necessities should be explained in terms of truths about essence; and they find the relevant essentialist truths in the normative case implausible. Consequently, they develop a *sui generis* modal notion of normative necessity to capture the truth in the vicinity of SUPERVENIENCE they believe to be consistent with most of our modal/ethical intuitions. Kramer (2009), a moral realist, rejects the conceptual truth of supervenience because he finds all the arguments for it unconvincing—e.g., those offered by Hare, Blackburn, and Shafer-Landau. Among other things, he thinks denials of supervenience are perfectly intelligible.

I ultimately reject the conceptualist line on supervenience not because I am an essentialist about metaphysical necessity or because I have strong pre-theoretic intelligibility intuitions either way. I reject it because I'm independently persuaded by an expressivist account of normative concepts. And I take myself to have shown that a motivation-expressive metasemantics together with constitutive facts about our motivations suggest that supervenience could not be conceptual truth. Moreover, I favor the normative

explanation at least in part because expressivists, unlike, say, synthetic naturalists, can preserve the *a priori* and obvious-seeming character of supervenience despite treating it as a normative truth.

[42] Contrast, e.g., the ethical naturalist's account of normative judgments as descriptive and as being made true by naturalistic or empirical truth-makers. It is not clear that the naturalist can guarantee the *a priori* nature of SUPERVENIENCE construed as a normative truth.

[43] Doesn't self-regard involve pure identity-dependence? It's not at all clear that it does. Our tendency to privilege our own pains and pleasures need not be construed in terms of a desire to privilege one's self on account of one's identity. It might instead involve a privileging of pains and pleasures one is intimately acquainted with.

[44] Since expressivists end up with a different take on supervenience than we might have started out with, here's how we should understand the dialectic. Supervenience seemed and has been taken by many (though not quite everyone) to be a conceptual truth. It was a reasonable starting point to take appearances at face value. But when we drew out the implications of an account of normative concepts that some of us find independently attractive, we discovered that supervenience couldn't be a conceptual truth if the account is true. This seemed *prima facie* to be a problem for the general account. But, fortunately, we discovered that the account can explain why supervenience might have seemed like a conceptual truth despite treating supervenience as an ordinary normative truth. It is consistent with the *a priori* character of supervenience, our willingness to reflectively endorse it, the fact that supervenience reflects a non-arbitrary constraint on normative theorizing, and so on.

Here's a general defense of proceeding this way. Except in the most obvious of cases (<every vixen is a female fox>), it seems reasonable to place less than maximal credence in whether a truth counts as conceptual or analytic. When the concepts in question are somewhat obscure, we should consider whether initial impressions of analyticity survive a full and comprehensive examination of the concepts. If they don't, we should still try to explain our initial impressions based on our theory of the concept. Otherwise, our theory is less than fully satisfying.

References

Atiq, E. (2016) "How to be Impartial as a Subjectivist." *Philosophical Studies* 173: 757–779.

———. Manuscript. "Ordinary Ethical Confidence."

Audi, R. (2015) "Intuition and Its Place in Ethics", *Journal of the American Philosophical Association* 1: 57–77.

Bader, R. (2016) "The Grounding Argument Against Non-reductive Moral Realism." *Oxford Studies in Metaethics*.

Blackburn, S. (1971) "Moral Realism," in *Essays on Quasi-realism*, Oxford: Oxford University Press.

———. (1984) *Spreading the Word*. Oxford: Oxford University Press.

———. (1993) "Supervenience Revisited," in *Essays on Quasi-realism*, Oxford: Oxford University Press.

Bratman, M. (1987) *Intentions Plans, & Reasons*. Center for the Study of Language and Information.

Daskal, S. (2009) "Plan-Based Expressivism and Innocent Mistakes," *Ethics* 119: 310–35.

Dancy, J. (1981) "On Moral Properties," *Mind*, 1981, 90: 367–385.

———. (2004) *Ethics Without Principles*, Oxford: Oxford University Press.

Dreier, J. (1992) "The Supervenience Argument Against Moral Realism," *Southern Journal of Philosophy*, Vol. 30.

———. (2006) "Negation for Expressivists: A Collection of Problems with a Suggestion for their Solution," in (ed. Shafer-Landau, R.) *Oxford Studies in Metaethics*, Vol. 1.

———. (2015) "Explaining the Quasi-Real," in (ed. Shafer-Landau, R.) *Oxford Studies in Metaethics*, Vol. 10. Oxford University Press. Oxford.

Enoch, D. (2011) *Taking Morality Seriously*. Oxford UK. Oxford University Press.

Fine, K. (2002) "Varieties of Necessity," in (Tamar Szabo Gendler & John Hawthorne, eds.) *Conceivability and Possibility*. Oxford University Press. pp. 253–281.

Gibbard, A. (1990) *Wise Choices, Apt Feelings*, Cambridge, Mass: Harvard University Press.

———. (2003) *Thinking How to Live*. Cambridge, MA. Harvard University Press.

Griffin, J. (1996) *Value Judgment*, Oxford: Clarendon Press.

Hawthorne, J. (2002) "Practical Realism?," *Philosophy and Phenomenological Research*, 64: 169–178.

Hazen, A. (2012) "Reflections on Counterpart Theory," in (eds. Bennett, K. & Zimmerman, D.W.) *Oxford Studies in Metaphysics* Vol. 7.

Heller, M. (1998) "Property Counterparts in Ersatz Worlds," *Journal of Philosophy*, 95: 293–316.

Jackson, F. (1998) *From Metaphysics to Ethics*, Oxford: Clarendon Press.

Jackson, F., Pettit, P, & Smith, M. (2000) "Ethical Particularism and Patterns", in *Moral Particularism*, Oxford: Oxford University Press.

Kim, J. (1984) "Concepts of Supervenience," *Philosophy and Phenomenological Research* 45: 153–76.

Klagge, J.C. (1988) "Supervenience: Ontological and Ascriptive," *Australasian Journal of Philosophy* 66: 461–70

Kramer, M. H. (2009) "Supervenience as an Ethical Phenomenon," in *Moral Realism as a Moral Doctrine*. Oxford: Wiley-Blackwell, 304–64.

Lewis, D. (1983) "New work for a theory of universals," *Australasian Journal of Philosophy*, 61: 343–377.

———. (1984) "Putnam's paradox," *Australasian Journal of Philosophy*, 62(3): 221–236.

———. (1986) *On the Plurality of Worlds*, Blackwell: Oxford.

McPherson, T. (2012) "Ethical Non-Naturalism and the Metaphysics of Supervenience." In *Oxford Studies in Metaethics* Vol 7.

Parfit, D. (2011) *On What Matters*, Oxford, UK: Oxford University Press.

Rosen, G. (forthcoming). "Normative Necessity." In M. Dumitru (ed.), *Metaphysics, Meaning and Modality: Themes from Kit Fine*. Oxford University Press.

Smith, M. (1994) *The Moral Problem*. Oxford: Blackwell.

———. (2004) "Does the Evaluative Supervene on the Natural?" reprinted in *Ethics and the A Priori*, Cambridge: Cambridge University Press.

Sturgeon, N. L. (2009) "Doubts about the Supervenience of the Evaluative," In *Oxford Studies in Metaethics*, Vol. 4.

Williamson, T. (2006) "Conceptual Truth," *Aristotelian Society Supplementary* Vol. 80: 1–41.

Zangwill, N. (1995) "Moral Supervenience," *Midwest Studies in Philosophy*, 20: 240–62.