

MORAL SUPERVENIENCE AND MORAL THINKING

Dalia Drai

ABSTRACT

The paper aims at meeting Blackburn's challenge (1971, 1984, 1985) to explain the non-reductive supervenience of moral predicates on natural ones. It offers a critical examination of Hare's model of moral thinking (1981) which can be used as a candidate for such an explanation. It is argued that, as it stands, Hare's model fails to meet Blackburn's challenge. Yet some revisions of the model are suggested, and it is claimed that the improved version does supply the required explanation. The model suggested in the paper carries a meta-ethical implication concerning the superiority of the anti-realist understanding of moral discourse.

Various philosophers in the field of ethics have claimed that ethical predicates supervene on natural ones. This claim has found consensus among philosophers with very different approaches to ethics such as Moore (1922) the intuitionist, Hare (1952) the prescriptivist, Mackie (1977) the subjectivist, Blackburn (1985) the projectivist, and Price (1986) the objectivist. Controversy does exist however, regarding the question of how to explain the phenomenon of supervenience, and particularly, what meta-ethical conclusions can be drawn from it.

In this paper I develop some of Hare's insights about moral thinking in order to meet Blackburn's challenge of explaining supervenience. In the first section I clarify the supervenience thesis in ethics and present Blackburn's challenge. In the second section I derive an explanation of supervenience from Hare's model of moral thinking, and show why this explanation does not fully meet Blackburn's challenge. In the third section I suggest a variation on Hare's model of moral thinking. This new model is a more faithful representation of moral discourse, and meets fully Blackburn's challenge. In the final section I discuss briefly the role of the anti-realist assumption in both Hare's explanation and mine.

SECTION ONE

The supervenience of ethical predicates on natural predicates is the thesis that if two actions (or two people) are indistinguishable with respect to their natural predicates, the same ethical predicates apply to them as well. This thesis should be understood as a modal one.¹ We do not only pronounce moral judgements on actions that have taken place, we also judge those that will take place and those that will not. When resolving an ethical dilemma, we do not only evaluate the action that we are going to take, but also the one that we have decided against taking. Furthermore, our moral judgements with respect to people do not only apply to actual people, but also to possible ones: we make judgements on people in counterfactual situations. Hare (1952, p. 145), for example, claims that not only do we judge Saint Francis as a good man, we also believe that if someone had the same natural qualities as Saint Francis, he would also be a good man. What are the meta-ethical assumptions underlying the discussion of supervenience in ethics, and what are the objections that can be raised regarding those assumptions?

The first assumption is that it is possible to speak about ethical predicates. The objection to this denies that there are ethical predicates at all, and claims that there are only ethical propositions. Ethical propositions are created by applying an ethical operator on non-ethical propositions. For example: "people take other peoples' lives" is not an ethical proposition, but the operator "it is wrong (that)" applied to that proposition creates the ethical proposition: "it is wrong that people take other peoples' lives".

The objection is ill grounded. Holding the supervenience thesis does not commit one to the superiority of a logical analysis of ethical propositions that includes ethical predicates, over analyses that do not. In order to be able to speak about supervenience, it is sufficient for there to be a *possibility* of analyzing ethical propositions as including ethical predicates, and indeed there is. To take the same example: the proposition "it is wrong to take another person's life", can be analyzed as the application of the ethical predicate "wrong" on every event in which someone takes another's life. Even if the objector insists that there are ethical propositions, but no ethical predicates, it is still possible to discuss the global supervenience of ethical propositions on non-ethical ones.² What I have to say about how to explain supervenience is equally relevant for the global version.

A further objection at this stage is that the distinction between ethical and non-ethical predicates is obscure, or even that different meta-ethical positions will yield different definitions of the distinction. However it is sufficient for the

¹ While it is important that supervenience has modal force, it is not relevant to the discussion in this paper to distinguish between the different kinds of modal force that we attach to the supervenience thesis. So for example I will not distinguish between weak supervenience and strong supervenience (Kim 1984).

² See Haugeland (1982) for a precise definition of global supervenience.

purpose of this paper that there exists a (perhaps open) list of predicates about which there is widespread agreement that they are ethical, even if there is no such agreement about a characterization that applies to all of them. The question about supervenience is the question about the relation between these and other predicates which are clearly non-ethical.

Questions can also arise about the set of base predicates. Does it include all non-ethical predicates, including for instance, aesthetic ones? Generally, writers on supervenience have taken the set of natural predicates as the set of base predicates. But this set is as difficult to characterize as is the set of ethical predicates. One widely accepted definition characterizes a natural predicate as one which can be discovered to be true of an object through empirical means. Sometimes a stronger definition is argued: a natural predicate is such that its application to an object can be verified by the methods of the natural sciences. It is sufficient for my purposes that there exists a wide list of predicates on which there is common agreement that they are natural.

The second assumption which underlies the discussion of supervenience in ethics is that ethical propositions can be true or false, and are not mere pseudo-propositions as Ayer argued (1936, p. 143). This assumption is necessary in order to make the claim that two objects whose natural descriptions are indistinguishable are also ethically indistinguishable. In order to say that they are ethically indistinguishable it is necessary to be able to say that ethical predicates can be true of them, and this is clearly at odds with Ayer's claim.

Let me anchor the second assumption in the meta-ethical debate on the descriptivism. Like Mackie (1977, p. 23-24), I take this debate to be about the meaning of ethical propositions, and more particularly, about the nature of the speech acts which are performed when making moral claims and uttering moral propositions. The descriptivist holds that the speech act performed in a moral claim is none other than the description of an object or objects, while the anti-descriptivist thinks that we are doing more: commanding, commending, expressing emotion, threatening and the like. Descriptivists are committed to the view that moral propositions, like all descriptive propositions, have truth values, and they are therefore in a position to discuss the question of supervenience, although not necessarily to affirm it. With anti-descriptivists the situation is more complex. The question of supervenience will be meaningless to any anti-descriptivist who gives up the notion of truth or correctness in ethics. It is meaningful, on the other hand, to an anti-descriptivist who believes that while ethical propositions are not descriptive they nevertheless do have truth values. To summarize, the second assumption necessary for any consideration of supervenience is the meta-ethical claim that the extreme anti-descriptivist is mistaken. However, this does not force a choice between either descriptivism or one of the moderate anti-descriptivist views.

Much of the current debate on supervenience is motivated by the challenge posed by Blackburn (1971, 1984, 1985) to explain why ethical predicates supervene on natural ones without being reducible to them. Blackburn

argues that the moral domain supervenes on the natural one in the following sense: it is necessary that if a particular object is B^* and A , then every object that is B^* is A . B^* is the complete characterization of the object from a naturalistic standpoint, and A is an ethical predicate (he calls this thesis S).³ The above necessity is a conceptual one: it is stronger than physical necessity, and according to Blackburn, it is also stronger than metaphysical necessity. Someone whose moral judgements are not subject to this constraint is not an immoral person; he simply does not understand what a moral value is, he has not mastered the moral language.

Blackburn distinguishes between the thesis that A supervenes on B , and the following thesis: every object that is B^* must necessarily be A (he calls this thesis N). This thesis guarantees a naturalist sufficient condition for any ascription of an ethical property, hence it is a claim about the reducibility of ethical predicates to natural ones. According to Blackburn, the thesis may be true without expressing a conceptual necessity: it is not a conceptual error to claim that Hitler was a good man; someone who makes this claim is flawed morally, rather than conceptually. The difference between S and N follows from the fact that it is conceptually true that if x and y are identical from a naturalist standpoint, they are also ethically identical. There is no conceptual truth, however, that determines that having certain naturalistic properties is a sufficient condition for a specific ethical property. Knowing that N is true does not only require conceptual capacity, it requires moral capacity.

Blackburn expresses the combination of S and the negation of N by recourse to conceptually possible worlds. There are possible worlds which contain objects which are B^* and A , and others which contain objects which are B^* and not A . But there is no possible world which contains objects of both types. Blackburn claims that although this partition of possible worlds is not a contradiction, it is mysterious and requires explanation.

This is the ban on mixed worlds: it is a ban on inter-world travel by things which are, individually, at home. the problem which I posed is that of finding out the authority behind this ban. (Blackburn, 1985, pp. 53-54).

The mystery is created by the fact that a conceptually possible world is one which can be conceived.⁴ But if we can conceive of a world with an object which is B^* and A and one in which an object is B^* and not A , why can we not conceive of a world which contains both these objects? This is Blackburn's challenge. Blackburn also claims that his challenge is more pressing for the

³ In the 1985 article, Blackburn offers a different definition of B^* : the set of natural properties that underlie property A . Since I do not believe the concept of "underlying" can be understood without supervenience, I will not use this characterization. To say that A supervenes on B is to clarify the intuition that B-properties underlie A-properties.

⁴ Conceiving is not exclusively a psychological concept; it is one which is subject to rational constraints.

moral realist. Whether this is so is a difficult issue which I will postpone to section four.

SECTION TWO

The fact that moral statements are not fully descriptive does not turn moral discourse into an exchange of subjective emotional attitudes outside the realm of reason.

Moral judgements about particular things are made for reasons, and the notion of a reason, as always, brings with it the notion of a rule which lays down that something is a reason for something else. (Hare, 1963, p. 21)

The above insight is central to Hare's model of moral thinking, and it serves here as the premise in what I will call "Hare's argument" for supervenience.⁵

Hare's argument is built on the claim that in order to justify the ascription of moral predicate *A* to an object, we must refer to another predicate *B* which is true of the object and to a rule relating *A* to *B*. If *B* is a moral predicate, then there must be a predicate *C* (different from *B*) true of the object, and a rule relating *C* to *B*. On the assumption that there is not an infinite or circular chain of justification, it follows that a non-moral predicate will eventually be arrived at. I can therefore assume that *B* is that non-moral predicate. Thus we arrive at the point that any ascription of ethical predicate *A* presupposes the existence of the rule $(x) (Bx \rightarrow Ax)$ where *B* is a non-ethical predicate. The derivation of supervenience proceeds as follows: assume that *x* and *y* have identical non-ethical properties, and that *x* is *A*. There is a rule justifying the relation between *A* and *x*, which we will assume to be $(x) (Bx \rightarrow Ax)$. *B* is a non-ethical predicate true of *x*, and therefore also true of *y*. From the rule it follows that *y* is also *A*, and that is the claim that ethical predicates supervene on non-ethical predicates.

Like every argument that derives supervenience from facts about justification, this derivation of supervenience must contain an anti-realist premise. The anti-realist premise is that truth (moral truth in our context) is an epistemic notion. Truth has to be defined similarly to the way Putnam (1981, pp. 49-74) and Dummett (1982, pp. 93-94) define it: a proposition is true if and only if it is justifiable in ideal epistemic conditions. When supervenience is based on justification, the distinction between two theses of supervenience is blurred. The first thesis of supervenience is a claim about our judgements: a set of predicates *A* is supervenient on a set of predicates *B* if for every *x* and for every *y* that are indistinguishable with respect to *B*, whenever we are justified to ascribe a predicate from *A* to *x* we are also justified to ascribe this

⁵ I am calling the argument Hare's argument because it relies on Hare's model of moral thinking. However, this paper should not be taken as a textual commentary on Hare's conception of supervenience.

predicate to y . The second thesis of supervenience is a claim about the objects themselves: a set of predicates A is supervenient on set B if every x and every y that are indistinguishable with respect to all B predicates, are also indistinguishable with respect to all A predicates.⁶

To understand the difference between the two theses of supervenience it will be helpful to examine the following example: B is a set of predicates describing the side on which a coin falls before time t . A is a set of predicates describing the side on which a coin falls after time t . Let us assume that a coin x was tossed 10^{10} times before t , each time landing heads, and that the same thing happened with another coin y . In such a case x and y are not distinguished with respect to B . We are thus justified in claiming that x will land on heads after t . For the same reasons we are justified in claiming that y will fall on heads after t . Yet it does not follow that if x falls on heads after t , so will y . The first thesis of supervenience applies here, but the second thesis, which deals with the question of how the coin will actually fall, does not. Hare's argument grounds the first thesis of supervenience but not the second. The only way to move from the first to the second thesis is by assuming anti-realism, because a realist has to respect the gap between what we are justified to say and what is really true.

This point is not intended as a criticism of Hare, since in the sense of anti-realism defined above, Hare is an anti-realist too. Hare (1963, p. 94) compares his doctrine to ethical theories which identify a true ethical judgment in a given situation with the judgment that would be made by an ideal observer in the same situation. There are disagreements about what to demand of the ideal observer. It is generally agreed that she would know all the relevant facts, understand all the moral terms involved, and have an appropriate capacity to use her imagination. But theories differ as to whether she should have a highly developed ability to empathize with the agent, or whether she should be ideally devoid of emotion. Whatever we decide on the specifications of the ideal observer, if we identify her judgement with truth, we are accepting anti-realism.

But I think that Hare is wrong to assume that every justification of an ascription of an ethical predicate requires the existence of a rule relating non-ethical predicates to an ethical predicate. I do not believe that the concept of justification necessitates a reference to rules. It is noteworthy that Hare himself claims that in order to justify a moral judgement one need not refer to rules, since it is sufficient simply to know that there are rules. But even in this weak sense it is incorrect to say that justification requires rules. We often justify the ascription of predicate A to x by referring to other properties of x , without believing that those properties are a sufficient condition for A . Dancy, in his criticism of the universalization requirement (Dancy, 1981, pp. 375-377), suggests the following example as a way of illustrating why the relation

⁶ Klagge (1988) describes this distinction, albeit for different reasons. He calls the first thesis Ascriptive Supervenience and the second Ontological Supervenience.

between moral justification and rules is far weaker than Hare supposes. A person ascribes to a certain girl the property "good" on the basis of her being chaste and pious. Therefore chastity and piety constitute reasons to ascribe the moral predicate "good" to her. But that does not mean that it would be justified to ascribe goodness to any girl who was chaste and pious; she might also be, for instance, cruel and therefore not merit the moral predicate of "good". This does not contradict, however, the fact that in the case of the first girl her chastity and piety did indeed constitute justifications for the judgement that she was good. Dancy summarizes his point in saying:

it is not the case that where my reasons for calling an action good are that it has properties ABC, I am committed to calling any other action which has the properties ABC good (Dancy, 1981, pp. 377).

I fully agree, and this is of crucial importance to the purpose of my argument, as Hare's rule-based view of justification makes his position especially vulnerable to Blackburn's challenge. Hare's assumption that every justification is based on a rule implies that in every case in which a moral predicate is ascribable to an object, there is a sufficient condition for the predicate, which is formulated in non-moral terms. This relation corresponds to Blackburn's *N*. Hence, Hare's argument does not meet Blackburn's challenge to explain supervenience without reduction.

SECTION THREE

Despite my objection to Hare's exaggerated use of rules, the general point of his argument is correct, namely that in certain cases, facts about justification entail supervenience. Until now I have presented only a small part of Hare's model of moral thinking, and a closer look at some of the details will serve as an introduction to my own model. Let us start with the description of a moral debate (although Hare's model is actually wider and can be applied to any moral thinking).

Suppose there is a person *A* who argues that in a given situation one ought to do *x*, while *B* maintains the opposite. According to the subjectivist the argument can end here, since *A* is simply expressing one kind of attitude toward *x*, and *B* another kind. Hare, however, argues that the demand for moral justification is always legitimate, and that at this point *A* is obligated to justify his judgement that *x* is the right action under the circumstances (and *B* of course is obligated to supply her own counter-justification). These justifications have a distinct form: *A* has to describe a natural aspect (*N1*) of *x*, and argue that action *x* is recommended by virtue of the fact that it is *N1*. Since there is a law concealed within every justification, *A* is committed to the generalization that in every situation identical to the given one with respect to its universal properties, an action having property *N1* is the right action.

In response, *B* may employ the following tactic: she can describe a situation identical with respect to its universal properties to the given situation, but in which the roles and characters are altered (for instance in the revised example, *A* may stand to suffer from the consequences of action *x*). *B* may then strongly suggest that *A* would not recommend performing the action that was *N1* in the imagined situation. It is this sort of tactic that Hare intends when he claims: "In this respect all moral arguments are *ad hominem*" (1963, p. 111).

Now *A* has two options before him. Either he can insist on the *N1* action even for the new situation described by *B*, in which case *B* failed to refute *A*'s argument, or he can admit that *N1* is not the right action for the new situation, but point to a relevant difference (or differences) between it and the original circumstances that justifies the difference in recommended actions. This difference in circumstances generates a difference between the description of action *x* in the original situation and the action resembling *x* in the new situation, which can be expressed by a new predicate *N2*, such that action *x* in the original situation was *N1* and *N2*, while the relevant action in the new situation is *N1* and $\sim N2$. Now *A* will give a revised account of the rule justifying action *x* in the original situation. Instead of being "All actions which are *N1* are justified in situations similar to the given one"; he will invoke the rule: "All actions which are *N1* and *N2* are justified in situations similar to the given one".

That will bring us back to the point at which the argument began. *B* now attempts to refute the new rule. It is worth noting that *B* would not make any headway in her effort to refute *A* by invoking a vastly different or strange situation in which an application of *A*'s rule would be undesirable in itself or produce undesirable consequences. For a legitimate refutation of the rule, *B* must describe a situation ostensibly "identical" to the original, but where *A* will agree that the rule does not apply. That is because in this sort of argument, what is sought is not a correct moral theory or an all embracing set of rules covering every possible situation. Instead, what is sought is a justification of a particular action in a particular situation. The appropriate rule will therefore be one which recommends a like action in a like situation (to the given one).

This move is sufficiently familiar from countless moral arguments that further examples will not be necessary (a detailed one is in Hare, 1963, pp. 90-111). For my purposes it is the formal structure which I would like to emphasize. Not only is the argument familiar, it is also legitimate and rational. A comparison to the philosophy of science can be of some help here in clarifying what Hare is after. Hare argues with some force (1963, p. 87), that moral arguments resemble scientific ones as depicted by Popper. An original hypothesis is proposed without justification; it is a mere guess. Science (and moral debate) is the attempt to refute the proposed hypothesis. I would like to stretch this analogy so that it also applies to the *philosopher* of science (and morals). The philosopher of science aims to uncover the structure of rational scientific debate by critically examining the activity of scientists. She thus

assumes that the scientist's concrete activities are relevant to her abstract model, even though the model would not necessarily be undermined by the fact that many scientists deviate from it. Moral philosophers similarly search for a model of rational moral debate. This model would not be invalidated by examples of people who deviated from it, unless it appeared that these deviations were themselves rational. Therefore the proposed model of the moral philosopher, as that of the philosopher of science, must accord with our notions of rationality, in particular with our judgements about the degree to which a specific moral or scientific argument is rational.

Following the above methodology of the moral philosopher, I want to suggest a modified version of Hare's model. Let us again suppose that *A* argues in a given situation that one ought to do *x*. I agree with Hare that the demand for justification is always legitimate, and that ultimately *A* will be compelled to justify his judgement with a claim of the form "*x* is *N1*", where *N1* is a non-ethical predicate. But I do not agree with Hare that *A* is thereby committed to some generalization. *B* will, in various ways, call *A*'s attention to other situations in which an *N1* action is not morally recommended. *A* can show *B* that in all those other situations the *N1* action is not *N2*, while in the given situation it is, and that what actually justifies the claim that *x* is the morally right action is the fact of its being both *N1* and *N2*, etc.

An example of such a procedure may be helpful. *A* argues that *B* ought to give *A* a thousand dollars. His justification is that a week earlier he lent *B* a thousand dollars (that is, *N1* = a week earlier *A* lent *B* a thousand dollars). *B* describes to *A* situations in which there is no obligation to return loans, for instance a case in which the lender plans to use the returned loan to cause harm to an innocent person (*N2*). *A* responds that the present situation is not like that, but is rather *N1* & \sim *N2*, and that what justifies his demand is that the situation is *N1* & \sim *N2*. *A* and *B* have made no reference at any point to generalizations. They may have referred to *prima facie* generalizations (Ross, 1932, Chapter 2 *passim*), but not to real ones lacking exceptions. A process like that will end when *A* or *B* are either convinced, or have given up trying to convince each other. But logically the process can continue indefinitely. An exceptionless generalization will never arise.

This model for the justification of ethical claims can also be applied, in my view, to the justification of aesthetic claims by means of non-aesthetic ones⁷ and the justification of counterfactuals by categorical propositions. Suppose that *A* argues that if I had come to a party I would have had a good time, while I claim that he is mistaken. In order to justify his claim, *A* refers to a categorical proposition: they played good music in the party. I in turn refer to a different categorical proposition: I had a headache. The claim that there

⁷ When we apply aesthetic terms, we justify applying them by referring to features which do not depend for their recognition upon an exercise of taste "... delicate because of its pastel shades and curving lines" or "... it lacks balance because one group of figures is so far off to the left and is so brightly illuminated" (Sibley 1959:424). When no explanation of this kind is offered, it is legitimate to ask or search for one.

was good music fails to justify the counterfactual; to do that one needs the conjunction of there having been good music and my not having had a headache. *A* may at this point agree that I would not have enjoyed myself, or he can refer to a different categorical proposition. The model is identical: in order to justify claims of a certain type, we refer to claims of another type. We have neither to refer to generalizations nor to assume that they exist.⁸

It is possible to deduce from my model of justification a supervenience thesis which does not presuppose moral rules. To demonstrate both the strengths and weaknesses of the model I will present it in its general form. There are domains of predicates, *A* (in this case the domain of moral predicates) and *B* (in this case the domain of natural predicates). It is known that when justifying a proposition from *A*,⁹ it is necessary ultimately to refer to propositions from *B*. I will prove that the domain *A* is supervenient on domain *B*. The first stage will be to prove global supervenience, after which I will adjust the premises so that they will entail local supervenience.

Proof of global supervenience by *reductio ad absurdum*: suppose that w_1 and w_2 are possible worlds which are indistinguishable in respect of propositions of *B*. Assume that in w_1 the proposition *a* from *A* is true, while in w_2 it is false. Now use the anti-realist assumption about *A*: if *a* from *A* is true, then in principle it is justifiable. To recall, justification is made by referring to propositions from *B*. Suppose that $b_1 \dots b_n$ from *B* justified the judgement that *a* was true in w_1 . w_1 and w_2 are indistinguishable in respect of propositions of *B* (given). Therefore $b_1 \dots b_n$ are also true in w_2 , and we are therefore justified in claiming that *a* is true in w_2 . That contradicts the premise that *a* is false in w_2 (a further use of an anti-realist premise).

But it can be claimed against the proof that although $b_1 \dots b_n$ justify *a* in w_1 they fail to do so in w_2 , because in w_2 there are other relations between propositions of *A* and propositions of *B*, and people can justify *A* propositions with other *B* propositions, and even with propositions not from *B*. This objection is misguided. When we say that the proposition “*a* is true in w_2 ” is justified, we do not mean by that that people in w_2 justify *a*. My premise was that *a* is justified by *B* propositions because of the meaning of *a*. That is how we use *A* propositions. When we justify *A* in w_2 , we continue to use the language as we are using it now; we do not begin to converse in the language of w_2 , Kripke clarifies this distinction in another context:

... when I say that a designator is rigid, and designates the same thing in all possible worlds, I mean that, as used in our language, it stands for that thing, when we talk about counterfactual situations. I don't mean, of course, that there might not be

⁸ This is precisely Levi's (1969:304) argument about statistical deductions. Statistical inferences are justified in virtue of being applications of inductive rules, and not because one of the premises is a generalization or a statistical law.

⁹ By a proposition from *A* I mean a proposition containing *A*-predicates. Generally that will be a proposition ascribing an *A*-predicate to an object.

counterfactual situations in which in the other possible worlds people actually spoke a different language. (Kripke, 1980, p. 77).

The proof of local supervenience is very similar to the proof of global supervenience. The premise about justification will be adjusted as follows: propositions ascribing property *A* to *x* are justified by means of propositions ascribing property *B* to *x*. With this premise it is possible to prove the local supervenience of *A* on *B*. Suppose that *x* and *y* are indistinguishable with respect to predicates of *B*, and let us assume that *x* is *A_i* in *w₁* and *y* is not *A_i* in *w₂*. I will again employ the anti-realist assumption. if *x* is *A_i* then it is possible in principle to justify the proposition that *x* is *A_i*. The justification will be made by ascribing *B*-predicates to *x*. Since it is possible to ascribe to *y* the same *B*-predicates, we are justified in ascribing the property *A_i* to *y* as well. That contradicts the premise that *y* is not *A_i*. I have thus deduced supervenience of the strongest kind from facts about justification. This deduction is made possible by the anti-realist premise that truth in ethics is equivalent to ideal justification.

To sum up: Hare's argument derives supervenience from a model of moral thinking which involves rules connecting every moral predicate to a combination of natural predicates. As a result, Hare could not explain the mystery: how can a set of predicates *B* determine the set of predicates *A* without there being any relationship between the individual *A* predicates and *B* predicates? My argument, on the other hand, derives supervenience without assuming specific relations between the individual predicates. Hence my explanation meets Blackburn's challenge while Hare's does not.

SECTION FOUR

The anti-realist assumption plays a crucial role in explaining supervenience. Indeed, the assumption is indispensable to the derivations presented in the previous section. Moreover, the assumption is indispensable in any derivation of supervenience that starts from facts about justification, otherwise, the gap between truth and ideal justification is not respected (see pages 7-9). This has important meta-ethical implications, it can be seen as a first step in an argument that deduces an anti-realist conception of truth in ethics via an inference to the best explanation. The realist is challenged to explain non-reductive supervenience. In contrast to the anti-realist he cannot base his explanation on facts about justification. Although I disagree with Blackburn that it is impossible for the realist to explain supervenience, I do believe that existing explanations that do not assume anti-realism are problematic.¹⁰ The explanation offered in this paper is the only one which meets Blackburn's challenge without distorting the nature of moral discourse. As long as the

¹⁰ Noonan (1987) and Forrest (1988) offer possible explanations of supervenience without assuming anti-realism. For a detailed criticism see Drai (1999, chap. 5).

DALIA DRAI

realist does not offer a better explanation than the one suggested here we are entitled to infer from supervenience an anti-realist conception of truth.

Dalia Drai (ddrai@bgumail.bgu.ac.il)
Dep. of Philosophy
Ben-Gurion University of the Negev
P.O.B 653, Beer-Sheva 84105, Israel

REFERENCES

- Ayer, A.J. (1936) *Language Truth and Logic*, Harmondsworth, Pelican Books.
- Blackburn, S. (1971) "Moral Realism", in J. Casey (ed.) *Morality and Moral Reasoning*, London, Methuen.
- Blackburn, S. (1984) *Spreading the Word*, Oxford, Clarendon Press.
- Blackburn, S. (1985) "Supervenience Revisited", in I. Hacking (ed.) *Exercises in Analysis*, Cambridge, Cambridge University Press.
- Dancy, J. (1981) "On Moral Properties", *Mind*, 90, pp. 367-385.
- Davidson, D. (1991) "Three Varieties of Knowledge", in A. Griffiths (ed.) *A.J. Ayer Memorial Essays*, Cambridge, Cambridge University Press.
- Drai, D. (1999) *Supervenience and Realism*, Hampshire, Ashgate.
- Dummett, M. (1981) "Realism", *Synthese*, 52, pp. 55-112.
- Forrest, P. (1988) "Supervenience: The Grand-Property Hypothesis", *Australasian Journal of Philosophy*, 66, pp. 1-12.
- Goodman, N. (1983) *Fact Fiction and Forecast*, Indianapolis, Bobbs-Merill (fourth edition).
- Hare, R.M. (1952) *The Language of Morals*, Oxford, Clarendon Press.
- Hare, R. M. (1963) *Freedom and Reason*, Oxford, Oxford University Press.
- Hare, R.M. (1981) *Moral Thinking*, Oxford, Clarendon Press.
- Hare, R.M. (1984) "Supervenience", *Aristotelian Society Supp.*, 58, pp. 1-16.
- Haugeland, J. (1982) "Weak Supervenience", *American Philosophical Quarterly*, 19, pp. 93-103.
- Kim, J. (1984) "Concepts of Supervenience", *Philosophy and Phenomenological Research*, 45, pp. 153-176.
- Klagge, J. (1988) "Supervenience: Ontological and ascriptive", *Australasian Journal of Philosophy*, 66, pp. 461-470.
- Kripke, S. (1980) *Naming and Necessity*, Oxford, Basil Blackwell.
- Levi, I. (1969) "Are Statistical Hypotheses Covering Laws?", *Synthese*, 20, pp. 297-307.
- Mackie, J. (1977) *Ethics: Inventing Right and Wrong*, Harmondsworth, Pelican Books.
- Moore, G.E. (1922) *Philosophical Studies*, New York, Harcourt, Brace & Co..
- Noonan, H. (1987) "Supervenience", *Philosophical Quarterly*, 37, pp. 78-85.
- Price, A.W. (1986) "Doubts about Projectivism", *Philosophy*, 61, pp. 215-226.
- Putnam, H. (1981) *Reason Truth and History*, Cambridge, Cambridge University Press.

- Ross, D. (1930) *The Right and the Good*, Oxford, Clarendon Press.
- Sibley, F. (1959) "Aesthetic Concepts", *Philosophical Review*, 68, pp. 421-450.
- Wittgenstein, L. (1958) *Philosophical Investigations*, Oxford, Basil Blackwell.