

A Contextualist Reply to the Direct Argument
Matthew H. Slater (matthew.slater@columbia.edu)
Department of Philosophy
Columbia University

[Draft, 2/26/04 — Forthcoming in *Philosophical Studies* — do not cite]

Abstract

The Direct Argument for the incompatibility of moral responsibility and determinism is designed to side-step complaints given by compatibilist critics against the so-called Transfer Argument. I argue that while it represents an improvement over the Transfer Argument, it loses some of its plausibility when we reflect on some metalogical issues about normal modal modeling and the semantics of natural language. More specifically, the crucial principle on which the Direct Argument depends appears doubtful where context plays a role in evaluation of normative claims.

1 Introduction

Grant that the semantics of our language is a contextual affair. Two tokens of the same sentence may differ in their truth value depending on their embedded context. Sometimes we make context explicit. More often it stows away. An exchange on an episode of “The Simpsons” comes to mind. Homer and Barney have applied for a unique spot on a NASA mission and after a series of rigorous physical and mental trials, the director addresses them to give his decision: “Well, Gentlemen,” he says, “You’ve completed the trials and done very well. In a way, you’re *both* winners. But, in a more *accurate* way, Barney’s the winner. . . .” Obviously there is no one concept¹ picked out by ‘winner’ in this story. The director calls our attention to this fact by suggesting a difference in context. There’s *some* context according to which both men could be reckoned as winners. Small consolation to Homer, for in *another* context — one which differs from the other perhaps by enjoying some closer relation to “the actual state of affairs”, or by being more demanding or strict in some other sense — one man can be properly considered a winner.

Nothing here is new. But it is sometimes forgotten. Sometimes the role of context is not

forgotten, but adduced in presumption that a certain context is to be considered the “favored” context. It is the director’s “more accurate sense” that carries the day so far as who gets the assignment. And the presumption is supposed to be that Homer is not *really* a winner. Asked, “who is *really* the winner?”, the director would respond univocally: Barney.

Sometimes the italicized ‘*really*’ — as in “are we *really* free?”, “are we ever *really* morally responsible?”, “do we *really* know anything?”, “are there *really* any ordinary things?” — is misused. Its ordinary function is often to raise standards. If ‘France is hexagonal’ is reckoned true according to some perspective (from very high up) and set of standards, asking (in a certain tone of voice, perhaps), “but is France *really* hexagonal?” will no doubt get us to revise our earlier reckoning. It does so by changing the context of the evaluation. But the propriety (and force) of asking if such and such is *really* the case depends on whether we think higher standards can or should rightly be invoked. *That*, of course, is often a hard question to answer — for what could determine what *standards* we ought to invoke seems to be a question of further standards (and whence *those* standards?). And what if we disagree about the appropriateness of certain standards? Very often, I think, that is that; no more can be said. But better, if we can recognize the nature of our dispute, we can make progress by looking at the dispute not only from the outside, but also from the *inside*.

It seems to me that these considerations offer a plausible way for resisting the “Direct Argument” for responsibility incompatibilism.² Whereas previous objections to the Direct Argument purporting to establish the *invalidity* of a crucial transfer inference appear unsuccessful, I shall offer weaker strategy I find under-explored and promising. Whether its ramifications are acceptable, however, is another question I shall have to leave unanswered here.

¹ No one non-context-dependent concept, I should say.

2.1 Logical Preliminaries

As both the Transfer Argument and Direct Argument are arguments in modal logic, it may be instructive to recall a few basic metalogical facts. Standard predicate or sentential logic is a calculus of truth — its axioms, rules of inference, equivalencies, &c. are accepted on the basis of their ability to preserve truth, whether we accept as valid the arguments they license. A *modal logic*, on the other hand, is a kind of logic that concerns one of the various “modes” of truth³ — the acceptability of its axioms and rules likewise depends on whether we are prepared to accept the arguments they license and whether we think that the rules properly capture the meaning of the operators and connectives involved.

Now recall that there is a class of modal logics called ‘normal systems’. Normal systems have in common a modeling strategy that greatly improves their suppleness over the Leibnizian possible worlds model. A “standard model” for a normal system adds an “accessibility relation” R between worlds (expressing something like relative possibility or relevance). Modalized sentences thus become world-relative, so for example we might propose that ‘ $\Box P$ ’ is true at world α just in case P is true at every world β such that $\alpha R \beta$. Varying the accessibility relation R in a model allows us to add different axioms to our system. But some axioms come for free. For example, all normal modal logics include the axiom

$$(K) \quad \Box(P \rightarrow Q) \rightarrow (\Box P \rightarrow \Box Q).$$

This axiom affords a sort of modalized version of modus ponens. Call it ‘(T \Box)’ for the “transfer of

² That is, the thesis that responsibility is incompatible with causal determinism. The Direct Argument is simply a version of van Inwagen’s Transfer Argument for freedom incompatibilism, and while I suspect that similar considerations offer a way of staving off freedom incompatibilism, I shall consider only a reply to the Direct Argument.

³ ‘Grass is green’ is true, but perhaps not *necessarily* true: there could have been a world with purple grass. In addition to the various logics concerning the alethic modalities (logics that result from adding various axioms), logics governing other modalities are possible. ‘People do not unjustifiably murder others’ is unfortunately false, but it *ought to be true*. Or maybe some fact is true, but it is not *known to be true*. By offering new interpretations of our modal operators and axioms to match, we can work out logics for deontic, doxastic, epistemic, and temporal modalities.

necessity”.

$$(T\Box) \quad \Box P, \Box(P \rightarrow Q) \models \Box Q$$

This follows directly from our modeling. Intuitively, if both ‘ $\Box P$ ’ and ‘ $\Box(P \rightarrow Q)$ ’ are true in our world, then by our semantics, ‘ P ’ and ‘ $P \rightarrow Q$ ’ will be true in all those R -related worlds (whichever they turn out to be) — and thus so will ‘ Q ’ be true in that set of world, rendering ‘ $\Box Q$ ’ true in ours. A version of this transfer principle is the crucial step in both the Transfer Argument and the Direct Argument, to which we now turn.

2.2 The Transfer Argument and the Direct Argument

Consider a sentential operator that expresses our “unfreedom” or powerlessness to alter the truth of whatever sentence it prefixes: call it ‘**PL**’ for “powerless”. For example, the sentence ‘**PL**(Peru is a democracy)’ is true if and only if I am/was powerless to substantially change the political structure of the Peruvian government; to make it false that Peru is a democracy. This sentence is true in part because of the *meaning* of ‘**PL**’ and certain facts about my political influence. In general, and for the purposes of the argument, we are to understand the meaning of ‘**PL**’ as ensuring that prefixing it to any true sentence describing a state of affairs prior to my birth (to be safe, we’ll make it *long before* my birth) will result in a true sentence. Note as well that ‘**PL**’ appears to function in very much in the same way as modal operators such as ‘it is necessarily true that’ or ‘it ought to be the case that’. Hence we propose the following rule of inference, the “Transfer of Powerlessness”:

$$(TPL) \quad \mathbf{PL}(P), \mathbf{PL}(P \rightarrow Q) \models \mathbf{PL}(Q)$$

Here comes the argument. Take some state of affairs S in the very distant past: owing to the fact that S was true long before my birth, it must be the case that I’m powerless to alter its truth:

$$(1) \quad \mathbf{PL}(S)$$

Now consider a statement of determinism⁴: that the present state of things follows from the state of affairs S and the laws of nature L (however interpreted). For simplicity, let's just consider an arbitrary specific consequence of determinism: some present (or near future) state of affairs T over which I ordinarily take myself to have power (say, my eating toast for breakfast tomorrow). Causal determinism states that T follows from S and L .

The question is whether I'm powerless to alter the truth of T . So far, we assume that I'm powerless to alter the truth of S . Plausibly, I'm also powerless to alter the truth of L and the fact that S and L together imply T . Thus we have:

$$(2) \quad \mathbf{PL}(L),$$

$$(3) \quad \mathbf{PL}(S \wedge L \rightarrow T),$$

which is equivalent, of course to:

$$(4) \quad \mathbf{PL}(S \rightarrow (L \rightarrow T)).$$

The transfer principle, (T \mathbf{PL}), which, applied to (1) and (4) gives us:

$$(5) \quad \mathbf{PL}(L \rightarrow T).$$

Applied again to (2) and (5) yields the unhappy conclusion, contra ordinary assumption, that I am powerless to alter the truth of T .

$$(6) \quad \mathbf{PL}(T)$$

The argument is general, since the choice of T was arbitrary. It seems, therefore, that I am powerless to alter any truth whatever: roughly, that free will is incompatible with determinism. From here it is a quick jump to the conclusion that morally responsibility is incompatible with determinism. For if I'm not free to bring it about that " T " is false, then it seems that I should not be held responsible for the fact that T .

The principle that responsibility requires freedom (or "alternate possibilities") has been challenged. But there is an argument exactly analogous to the forgoing — the 'direct argument' — that

⁴ I'll necessarily leave the definition rather loose — I expect the point is clear enough.

argues for responsibility-incompatibilism without recourse to such a principle. We get it by swapping the ‘PL’ operator in (1–6) for one which expresses non-responsibility, ‘N’ (where ‘N(P)’ means something like: “P, and no human being, or group of human beings, is even partly responsible for the fact that P.”⁵) and offering a “transfer principle” analogous to (T□) and (TPL):

$$(TN) \quad N(P), N(P \rightarrow Q) \models N(Q)$$

Again, the conclusion that N(T), that no human is even partly responsible for T, follows for any truth T on the assumption that determinism is true. And nowhere in this second argument do we appeal to the premises that free will is incompatible with determinism or that our enjoying moral responsibility entails that we are free.

A few things to notice about these arguments right away. First, though quite similar, they are distinct arguments. If Frankfurt and friends are right that moral responsibility does not require “alternate possibilities”, then the transfer of powerlessness argument might be sound and the transfer of non-responsibility argument (the Direct Argument) invalid. Second, the validity of each argument turns most centrally on the validity of the relevant transfer principles, (TPL) and (TN).⁶ Otherwise, the arguments’ conclusion follow from their premises just in virtue of standard sentential logic.

So what can be said in defense these principles? As I mentioned above, — and as van Inwagen points out — any normal modal semantics will validate either transfer principle. No restriction on the accessibility relation R will result in a model on which (TN) is invalid. The question is whether we should give a normal modal semantics for ‘N’. Van Inwagen remarks that ‘N’ is “very like a modal operator”⁷ — but it is not clear on what basis this claim is made. There is a *prima facie* functional similarity between ‘N’ and other modal operators: it is a non-truth-functional sentential operator which attaches to declarative sentences. The central question, however, is whether it should be

⁵ Van Inwagen (1980), p. 32.

⁶ Though, of course, we might also challenge the other premises, this strategy would take us too far afield into a discussion of laws and determinism.

⁷ Van Inwagen (1980), p. 33.

modeled as a “normal” modal operator. The analogy is not enough. Van Inwagen concedes this point:

While this formal result is not devoid of persuasive force . . . it is far from decisive. It depends on the assumption that there is *some* set of worlds W such that $N(P)$ *can* plausibly be thought of as making the assertion that P is true in every member of W . While this assumption seems right to me, I have no argument for it, and a person who was determined to reject (TN) might very well reject it.⁸

While some philosophers apparently seem to accept the standard modal interpretation of ‘ N ’ as a truth of logic, it is important to see that this is a normative, not a logical question. It is an axiom schema of *a* logic, indeed, a large class of logics that have (unfortunately, I think) been dubbed ‘normal’. The normative question concerns the appropriateness of this logic to its intended subject matter. Insofar as ‘ $N(P)$ ’ is supposed to mean “ P and no one is or has ever been even partly responsible for the fact that P ”, whether it befits a normal modal semantics depends on how we reason about such claims.

There are other cases in which normal axiom schemas look problematic. Deontic logics which construe a ‘ OP ’ as ‘obligatorily P ’ might include the theorem:

$$(OD) \quad \neg(OP \wedge O\neg P)$$

But as Chellas points out,

it is a matter of controversy whether deontic logic should thus rule on the question whether contrary obligations OP and $O\neg P$ are always inconsistent. Indeed, it is arguable that the possibility of such conflict is a main feature of some concepts of obligation, that it is often this, for example that gives moral dilemmas their poignancy.⁹

Less controversial is:

$$(OD^*) \quad \neg O\perp,$$

in effect, a construal of ‘ought implies can’. The problem is that in any normal modal interpretation of ‘ O ’, (OD) and (OD*) are equivalent. Thus, the controversial (OD) and uncontroversial (OD*) are regarded as equivalent — that is, ‘ $\neg O\perp \leftrightarrow \neg(OP \wedge O\neg P)$ ’ is a theorem on any normal interpretation of

⁸ Van Inwagen (1980), p. 34. In van Inwagen’s paper, the principle I am calling ‘(TN)’ is labeled ‘(B)’.

‘O’. Regardless of what you think of either principle, we should reject an deontic logic that affords us “no way to distinguish the principle that *ought* implies *can* from the principle that conflicts of obligation cannot exist.”¹⁰

This objection does not of course directly call (TN) into question, but it makes clear that normal semantics sometimes fails as an account of the modal character of obligation. In the case of deontic logic, we need a weaker system that doesn’t blur distinctions we are interested in making. If such distinctions are controversial, however, other considerations must be brought to bear on the question of the appropriateness of the logic. The important thing to see is just that this question must come up — the presumption should not be on the side of the logic simply because it is *logic*. How, then, should we proceed to determine the validity of (TN)? Van Inwagen suggests that we should carefully consider instances of it in order to appreciate its plausibility. For example:

(1) N(John was bitten by a cobra on his thirteenth birthday)

(2) N(John was bitten by a cobra on his thirteenth birthday →
John died on his thirteenth birthday)

hence,

(3) N(John died on his thirteenth birthday)¹¹

As quotidian examples like this illustrate, the transfer principle indeed looks quite plausible.

2.3 Previous Responses

Let us grant for now that the validity of the Direct Argument turns on (TN). And let us also grant the premises that I’m responsible neither for the goings on in the distant past nor for the truth of determinism. To show that the transfer principle is invalid requires producing a counterexample of a certain special kind. Van Inwagen issues this invitation: we must attempt to construct a counterexample that is both *more plausible* than the transfer principle, and which does not rely on the assumption that responsibility compatibilism is true.¹² Traditional responses to the Direct Argument typically take up this invitation.

This is, for example, the central project of Ravizza’s (1994)-cum-chapter six of *Responsibility*

⁹ Chellas (1980), p. 201.

¹⁰ Chellas (1980), p. 202.

¹¹ Van Inwagen (1983), p. 187.

and Control. Fischer and Ravizza grant the plausibility of (TN) — that it often accords very well with how we reason about responsibility. But never mind: just one counterexample (of the invited type), however strange, will do for invalidating the principle, however natural.

Their centerpiece counterexample is “Erosion”: Secret agent Betty is set to the task of destroying an enemy camp by placing explosive charges at sensitive points of a mountain glacier in order to cause an avalanche. But little does anyone know that the glacier has been eroding and will soon avalanche. The scenario is a meat-and-potatoes “Frankfurt counterexample”. Betty detonates the charges at T_1 in order to crush the base at a crucial later time T_3 . Had she not, the glacier would have broken off suddenly at T_2 due only to the ravages of time, and resulted in the crushing of the camp at T_3 . Betty’s deeds preempted the natural process (for which no one at all is responsible) which would have resulted in the same conclusion. Is Betty responsible for the camp’s being crushed? Fischer and Ravizza want us to say “Yes!”, for Betty is in our ordinary estimation a paradigm case of a responsible agent &c. If so, we have

(1) N(The glacier is eroding)

(2) N(The glacier is eroding \rightarrow there is an avalanche that crushes
the enemy base at T_3)

but,

(3) \neg N(there is an avalanche that crushes the enemy base at T_3)¹³

And thus, that (TN) is invalid.

Such putative counterexamples have met with considerable, though often inconclusive, resistance. McKenna, for example suggests that the incompatibilist might claim that such examples are irrelevant to the question of the incompatibility of responsibility with determinism, for determinism concerns the *actual* scenario. He writes:

Determinism is not a theory about what happens in some counterfactual situation. The alleged

¹² Ibid., p. 187.

¹³ See Fischer and Ravizza (1998), p. 156–8.

threat of determinism concerns what transpires in the actual world. For the counterexamples to answer a properly modified version of (TN) . . . they must show that the Ensuring condition arises in the actual world.¹⁴

The properly modified version, McKenna argues, must take stock of the actual sequence of events and P's actual causal sufficiency for bringing it about that Q. He notes that the Fischer/Ravizza counterexample relies on preemptive overdetermination — or so-called “two-path cases” — whereas the incompatibilist need only argue from determinism in “one-path” cases, cases in which Q has only one sufficient cause. And it is impossible to come up with suitable counterexamples to such a revised principle. This we might grant — though I worry that it is asking too much.

The compatibilist's counterexamples seem to attain most of their force from the intuition that the counterfactual intervener has no effect on the situation. The idea is that the actual sequence of events in a Frankfurt scenario is identical to one in which the agent would be judged to be guilty/responsible. How can a merely *possibly intervening* event have any effect on my abilities? Peter van Inwagen replies convincingly:

Well, in a way it cannot — [the intervener's presence] cannot diminish one's skill as a marksman, or make one any less a master of disguise . . . but it *can* . . . affect one's abilities with respect to determining the truth-values of various propositions.¹⁵

Which are the propositions that matter? Suppose (to use van Inwagen's 1997 example) that both Sandy and Poisson want Jane dead. So before Jane sets out for a long trek in the desert, Poisson poisons her water bottle. And, not knowing what Poisson had done, Sandy empties it and fills it with Sand. Jane dies several days later (on the 12th of July, say). Question one: is Sandy responsible for Jane's death? This question is logically distinct from question two: is Sandy responsible for the fact that Jane died? The court laughs out Sandy's defense that he actually *extended* Jane's life a bit and answers the former question unequivocally: Sandy caused Jane's death and is thus responsible for the death that

¹⁴ McKenna (2001), p. 45.

¹⁵ Van Inwagen (1997), p. 377–8.

Jane actually died. But we may give such an answer to the first question while answering the second in the negative.

It seems to me evident that Sandy was not morally responsible for Jane's having died in the desert on or about 12 July. But he *was* morally responsible for her death. . . . And he could not have been morally responsible for her death if he had not been morally responsible for some of the facts relating to her death. . . . If, therefore, one decides on general philosophical grounds that Sandy was unable to act otherwise than he did — courts are not philosophical seminars; *counts simply take it for granted that people are in general able to act otherwise*, just as they simply take it for granted that sense-perception is in general reliable — then one should conclude that he was morally responsible for no facts relating to the case . . . and therefore not morally responsible for Jane's death.¹⁶

The standards we bring to the table in deciding how to evaluate such scenarios may well be the court's — but then we should be explicit about them: if the reason we find Sandy responsible for Jane's death (though perhaps not the truth of the proposition that Jane died in the desert) is that we assume that Sandy is responsible for filling the water bottle with sand, and thus responsible for the fact that Jane died of dehydration. That's fine, but we must be fair in our application and *remember* this assumption when we go to ask about other cases. That is, we cannot necessarily take the result here (that Sandy is responsible for Jane's death, despite the fact that Jane would have died no matter what) and go on to argue in general that alternate possibilities don't matter.

I don't want to go over too much old ground. I merely want to call attention — as before — to the dialectics of the situation. The transfer principle cannot be proven to be valid (everyone acknowledges), but it seems in many circumstances quite plausible. But of course, it is also quite plausible that Sandy is responsible for Jane's death and Betty is responsible for crushing the enemy camp. But neither are those cases very conclusive. Is this a stalemate?

2.4 *The Epistemic Parallel*

Before turning in earnest to a contextualist response to the Direct Argument, I'd like to consider

¹⁶ *Ibid.*, p. 379; my emphasis.

briefly a parallel dialectic. Many theories of epistemology include a very plausible transfer principle for knowledge, structurally exactly alike that considered above, where 'K(P)' is understood as "S knows that P":

$$(TK) \quad K(P), K(P \rightarrow Q) \models K(Q)$$

This principle is equivalent to:

$$(TK') \quad \neg K(Q), K(P \rightarrow Q) \models \neg K(P)$$

Now let 'P' be "Here are two hands before me" and 'Q' be "I'm not a brain in a vat" (or some other suitably compelling skeptical hypothesis). We then get a standard Cartesian skeptical argument.

- | | | |
|-----|---|----------------------|
| (1) | I don't know that I'm not a brain in a vat | $\neg K(Q)$ |
| and | (2) I know that if here are two hands before me, then I'm not a brain in a vat, | $K(P \rightarrow Q)$ |
| so, | (3) I don't know that here are two hands before me | $\neg K(P)$ |

Perhaps the most famous objection to this kind of transfer argument for knowledge is Moore's. *I do know that here is one hand and here is another. Any argument that suggests otherwise must be wrong, for my confidence in such pieces of knowledge outstrips the ability of any philosophical argument to show me otherwise. Philosophy is nothing but the sheerest speculation!*¹⁷ My hands are *right here*.

Many philosophers seem to think that Moore straightforwardly begged the question or misunderstood the skeptical challenge. But admittedly, there is something compelling in his response, if only because we should like to hold fast to our knowledge of the external world in the face of philosophical challenges. Similarly for our claims of responsibility. Might we then respond to the direct argument in the way that Moore responded to the skeptical argument?

The principle (TK), also known as the closure of knowledge under known entailment (or

“closure” for short), is quite plausible. We appreciate this fact in much the same way as we appreciate the plausibility of (TN): by reflecting on ordinary instances of it. Nevertheless, the closure of knowledge (along with some eminently plausible assumptions) seems to have the untoward consequence that we know next to nothing about the world. Yet, of course, that is false: we know all kinds of things.

Slote (1982) takes up the question of whether ‘N’ is a normal modal operator (and thus whether the transfer principles are valid) by picking up on the parallel between the transfer of knowledge and the transfer of non-responsibility. He uses the parallel to call into question the basis on which we saw van Inwagen argue for the standard interpretation of ‘N’: there are other sentential operators which look very much like they should be normal modal operators — ‘K’, and ‘O’ for example — yet which do not admit of a normal semantics, for our ordinary use of ‘K’ invalidates transfer principles which those semantics must validate.

But, Fischer asks, how far should we be moved by this analogy? On what basis, even, should we even accept that epistemic closure fails? Slote chalks this failure up to the apparent fact that “people may fail to make inferences they are entitled to make”. That’s a little vague. But we may also argue against the closure of knowledge directly from skeptical arguments employing it. Fischer cites Dretske’s “Zebra Example”¹⁸:

- (1) A knows that there is zebra in front of him,
 - (2) A knows that if there is a zebra in front of him, it is not a cleverly disguised mule,
- but, (3) A *does not* know that it is not a cleverly disguised mule in front of him.

¹⁷ David Lewis (1991) made a surprisingly similar plea regarding the duel between mathematicians and those (presumptuous!) philosophers who dared suggest that math is a subject with no object (or worse, false on “philosophical grounds”).

¹⁸ Fischer (1996), p. 33.

So rather than arguing back from (3) via closure [via (TK')] that (1) is false, Dretske asks us to deny (TK'). But he offers more defense than just the Moorean denial: Knowledge very often “consists in a certain sort of discriminatory capacity — the capacity to distinguish the actual situation from a range of relevant alternative scenarios”¹⁹ The reason both (1) and (3) can be true is because the scenario in which the animal *A* sees is a cleverly disguised mule is not *relevant* to *A*'s claim to know that there is a zebra in front of him. Again, there is something compelling about such a position, but it is awfully tenuous in the context of the skeptical argument. Fischer puts his finger on the crucial point:

It seems to be difficult to establish such a theoretical approach in light of the existence of intuitions highlighted by the skeptic. So, for example, how could one *argue* that *A* knows that there is a zebra in front of him?²⁰

Engaging in such an argument seems to promote irrelevant alternatives to relevancy. Fischer doesn't find the situation very promising for resolving the dispute between Dretske and the skeptic. But of the strategy Fischer finds “more promising for casting doubt on closure” — having to do with one's beliefs being “integrated” —, he doubts whether Slote will be able to find any interesting analogous strategy for questioning responsibility transfer.

I think we may most plausibly read Slote as issuing a challenge (along the lines of the considerations I raised above) to those like van Inwagen who are moved by the outward “modal character” of the relevant operators. The analogy is not supposed to work all the way down to an analogous diagnosis of why transfer fails in the responsibility case, but rather point out that it is often the case that operators which look to be normal modal operators can very plausibly turn out *not* to be. Of course, we would like to be able to say more in the responsibility case in the way of *explicit* reason why *these* operators do not admit of standard modal semantics.

¹⁹ *Ibid.*, p. 34.

²⁰ *Ibid.*

3.1 Moral and Epistemic Contextualism

I said in the introduction that we evaluate propositions on the bases of standards and that some standards appropriate in one context are not appropriate in another. The standards picture of semantics is of course an idealization — and good thing it is: as it encourages a hard-to-sustain hierarchical picture; of standards for standards, on and on, all the way up — when what we actually do is far more complex than simply applying some explicit criteria or other. But it is nevertheless illuminating reflect on the dynamics of standards so far as we can make them out.

When I claim to know that the Yankees lost the 2001 world series, I mean my claim to be favorably evaluated according to a certain set of standards. Perhaps the standards are to include certain possible alternatives from which I must distinguish actual events and exclude others: e.g., that they really won, but that I was duped by a huge conspiracy. Any standard that requires ruling out such an outlandish alternative possibility is quite exacting. But my claim to knowledge was probably not made in what I took to be an exacting context. Other statements in other situations — e.g., “the patient will die unless we amputate”, said in the operating room — may invoke *very* demanding standards indeed.

I think that something like this point is crucial to the proper diagnosis of the skeptical argument. I believe, as David Lewis put it, that ascriptions of knowledge are subtly context-dependent, and that epistemology is a context that makes them go false.²¹ That makes for the ironic claim that epistemology is a subject that “destroys its own subject matter”.²² The “diagnosis” thus allows that the skeptical argument is fine, as it goes, when we’re doing epistemology — but that it doesn’t in general sully claims to knowledge made when we’re *not*.²³

Now obviously knowledge has much to do with ordinary estimations of responsibility. I’m to a

²¹ Lewis (1996), p. 551.

²² *Ibid.*

²³ There are, I think, many implausible aspects of Lewis’s contextualism, but the general strategy seems worthy of pursuit.

great extent off the hook (considerations of moral luck or negligence aside) if I *unknowingly* back over my neighbor's kitten in my car. I'm blameworthy, on the other hand, if I *did* know that it was sleeping under my tire and backed up anyway. So what's to stop the skeptic from pressing the consequences of his argument to the moral realm? According to him, I *never* know whether there's a kitten under my tire. A number of responses are possible for those who take this crude challenge seriously — shifting the focus from knowledge to mere belief, citing facts about my intent &c. — but they miss the more important point: the “kind” of knowledge required for moral responsibility falls short of the requirements assumed by the skeptic. The discussion of my knowledge (and responsibility) takes place in my driveway, at the scene of the crime — not in “the philosophy room”. Charges are brought up in the courtroom where evidentiary standards are supposed to exceed *reasonable* — not philosophical — doubt.

This point may seem naïve, parochial, or even anti-philosophical, but I think that is substantially right. I may know that there's a kitten behind my tire when I'm looking at it in my driveway, but fail to know it when I'm reflecting philosophically on skeptical alternatives. I'm not exactly sure how to argue for this claim, though. In any case, there's no *conclusive* argument for it (if there's a conclusive argument for *anything*) because giving such arguments often require invoking the kind of high standards which a contextual theory must avoid in order to maintain the ordinary judgments it licenses.

If knowledge *is* contextual in the way I've suggested, what of the closure principle? I think we may say one of two things: one, that closure is false because it has counterexamples (say, of the sort that Dretske adduces); or two, that closure is correct so long as we hold context fixed. Subject to a certain standard (a certain epistemic context), the closure principle is valid. That is essentially to claim that the fallacy in arguing for the incompatibility of knowledge with closure (to put the skeptical point another way) is that the ‘K’ operator equivocates between different (perhaps “implicitly subscripted”) operators. In the “Zebras” example, we might have

(since I cannot relinquish my reactive attitudes), then I'm stuck failing to believe what's nevertheless *true*. My moral practices (however entrenched) belie moral *truth*.

Perhaps, at best, the practices pave the way for a suitable proxy notion, even if the genuine article is out of reach. Or perhaps the practices may themselves be justified instrumentally — at maintaining social order or what have you (deterministic world or no). These avenues, however, are retreats: we should want more. The Strawsonian may attempt to get a foothold by trying to secure a more modest, non-metaphysical moral objectivity. Does moral objectivity demand, as the hard-core realist presumes, non-natural moral facts or obscure and causally remote platonic forms? Answering in the affirmative makes the moral false consciousness of the hard-determinist understandable: for if moral facts *do* exist in Plato's heaven (alongside the numbers, I suppose), it is obscure how we could know anything about them, save by some queer faculty of non-natural intuition.²⁶

This picture of moral objectivity should be denied. But what's to come in its place? The question is notoriously difficult — though I'll have something to say presently about its prospects or how it might be filled out. We don't have to buy Strawson's line in detail to appreciate its fundamental point: moral categories are learned, and very often the justification for their application cannot be distanced from their learning. We learn how to apply moral categories (indeed, I think, *any* categories) by looking at how the terms are used — that is, by looking at moral *practice*.

Practice and Fact, however, are not disconnected notions. Navigational practices changed once we learned that the earth wasn't flat. Practice is rarely immutable. What shall we say to those critics of Strawson²⁷ who attempt to call into question his impossibility claim by pointing out that our moral practices *have* changed significantly even in the last two hundred years or so? What we make of such inductive evidence depends on whether we can draw a line in the sand between those attitudes which can be properly regarded as "immutable" (or whatever is needed) and those which are subject to

²⁶ such was the core of Mackie's arguments against moral objectivity.

²⁷ e.g., Benson

change. Whence our certainty that if only we were smart enough or met an appropriately convincing rational argument to the effect that we should give up our moral practices that we wouldn't snap out of our delusion and drop them at once? Rational argument has done much to controvert what must have seemed like incontrovertible intuitions.²⁸

This may be so. Though I find it quite unlikely that we will be able to completely abandon the reactive attitudes that Strawson thinks underlie ascriptions of moral responsibility, I do not think that we should base any argument on the premise that we *cannot*. Instead, I think that we should say something like the following: our practices of blaming and holding people accountable *express* certain standards, standards which are not beholden to metaphysical issues like determinism; and all that is needed for a kind of objectivity (corresponding to the naïve thought that the standards get things “right”, that they can be recommended to others) is for us to be prepared to *defend* these standards in a particular way and for them to fit in with our practices. Timothy Scanlon makes a provocative plea for this conclusion:

It is not necessary, in order to explain the possibility of being mistaken, to construe the relevant subject matter in a metaphysical way as existing outside us. The question of whether there are standards of the required sort is a substantive one within the subject in question — a matter of whether there are conclusions and ways of arriving at them that we have no reason to regard as defective. It need not be a metaphysical question about what exists or an epistemological one about how we are in touch with it.²⁹

Does the Direct Argument give us good reason to abandon these ordinary standards? That depends, perhaps, on our *goals* in doing philosophy. And of course, much more needs to be said about the nature of any standards and of their accompanying claims to normative authority.³⁰

Scanlon's rendering of moral objectivity at once frees us from having to posit awkward moral facts or provide external philosophical justification and paves the way for a contextualist account of

²⁸ I find the cases of special and general relativity especially perspicuous.

²⁹ Scanlon (1998), p. 63.

³⁰ on this count, I find Gibbard's (1990) *Wise Choices, Apt Feelings*, chapters 8–11, particularly illuminating.

responsibility. Different standards in different contexts do not reveal morality to be devoid of content any more than the existence of contextual standards for knowledge reveal epistemology to be devoid of content. It is no strike against normative discourse that it is compartmentalized. In one sense, on one set of standards, I am not responsible for a child starving in Africa. In another more stringent sense, considering how easy it would be for me to send a small (easily affordable) amount of money to UNICEF, and considering that I *don't*, it seems as though I *am* responsible for her death.³¹

Even if one accepts this sort of picture for epistemology³², one might worry about a disanalogy between contextualized epistemology and ethics.³³ The epistemologist can grant the contextual nature of knowledge without giving up on epistemology as a serious pursuit by “naturalizing it” and investigating its ties to neighboring notions of belief and justification (both objective notions, in their way) and setting aside evaluative questions such as “how much justification is required for knowledge”. Is a similar avenue open to the ethicist? How can she answer ordinary ethical questions that apparently depend on settling the question of compatibilism? Doesn't the “philosophical context” trump the “ordinary context”? This seems to me a pressing concern — one that is exceedingly hard to answer. First, insisting that there is no preferred context when it comes to answering ethical questions (even if that insistence issues from going ethical concerns) does not substantially differ from insisting that one context *is* preferred. Second, if there are numerous ethical contexts prescribing different courses of action, how is possible that there is any ethical truth?

Perhaps we can answer as follows: insofar as we find this latter *meta-ethical question* to be less pressing than *ethical questions* (such as determining whether I am responsible for backing over my neighbor's kitten), we may be entitled to leave aside the question of truth, perhaps in favor of an “operational” sense of objectivity. Perhaps all we do is start from within some context and work our way out, recommending outsiders join us. That is the central difference between subjectivism and

³¹ See Unger (1995).

³² And it is far from clear that one *should*.

(some brand of) objectivism: the claim, from within a context, that these are the standards that you should come to hold as well.

3.3 Conclusions

If we grant that our practices are at least in part constitutive of the meaning of ‘responsible’, and that our practices are governed by different standards of varying degrees of stringency, we are well on our way to denying the crucial transfer principle (TN). And even if we cannot straightforwardly *deny* it, we may shift the burden to its defenders to establish that the crucial operator is governed by a normal modal semantics. I have not yet said enough to show that claims of responsibility are contextual nor that a contextualized ethics would not collapse into subjectivism. A contextual theory of responsibility looks to me attractive both because it accords with our moral practices (the “semantic phenomenology” of moral claims, if you like) and because it affords a diagnosis of the Direct Argument. These are accomplishments for a philosophical theory for those who see philosophy as an extension of ordinary reason and practice. Those, on the other hand, who see philosophy as a revolutionary endeavor, may not be so impressed with such accomplishments. *Why should we think that our practices are reasonable or truth-resembling? If the Direct Argument is right and no one is ever responsible (supposing determinism is true), then those practices belie the truth.*³⁴ I don’t know what to say against such a position except perhaps just to hold on to the claim that we have good reason to think that our practices are *not* wrong-headed. Anyway, whatever your view of philosophy, I hope you agree that the Direct Argument does not look all that direct.

³³ This worry, such as I understand it in outline, was suggested to me by an anonymous referee.

³⁴ Richard Double in his very interesting (1996) discusses these issues in the context of different “metaphilosophies”. He calls the former goal (or something like it) ‘philosophy as underpinning common sense’ and the latter, a species of ‘philosophy as world-view construction’ (see pp. 22–30).

REFERENCES

- Chellas, B. (1980) *Modal Logic*. Cambridge University Press.
- Double, R. (1996) *Metaphilosophy and Free Will*. Oxford University Press.
- Fischer, J. M. (1996) *The Metaphysics of Free Will*. Blackwell.
- Fischer, J. M. and Mark Ravizza. (1998) *Responsibility and Control*.
- Gibbard, A. (1990) *Wise Choices, Apt Feelings*. Harvard University Press.
- Lewis, D. (1996) "Elusive Knowledge" *Australasian Journal of Philosophy* 74, pp. 549–567.
- McKenna, M. (2001) "Source Incompatibilism, Ultimacy, and the Transfer of Non-Responsibility" *American Philosophical Quarterly* 38.1, pp. 37–51.
- Ravizza, M. (1994) "Semi-Compatibilism and the Transfer of Non-Responsibility" *Philosophical Studies* 75, pp. 61–93.
- Scanlon, T. (1998) *What We Owe to Each Other*. Harvard University Press.
- Slote, M. (1982) "Selective Necessity and the Free Will Problem" *Journal of Philosophy* 79:1, pp. 5–24.
- Strawson, P. (1962) "Freedom and Resentment", reprinted in Watson (1982)
- Unger, P. (1995) "Contextual Analysis in Ethics" *Philosophical and Phenomenological Research* 55.
- Van Inwagen, P. (1980) "The Incompatibility of Responsibility and Determinism" *Studies in Applied Philosophy*.
- Van Inwagen, P. (1983) *An Essay on Free Will*. Oxford University Press.
- Van Inwagen, P. (1997) "Fischer on Moral Responsibility" *The Philosophical Quarterly* 47:188, pp. 373–381.
- Watson, G., Ed. (1982) *Free Will*. Oxford University Press.
- Widerker, D. (2002) "Farewell to the Direct Argument" *Journal of Philosophy* 99(6), pp. 316–324.