

More Bad News for the Logical Autonomy of Ethics

MARK T. NELSON
Westmont College
Santa Barbara, CA 93108
USA

I

Since the time of Hume, many philosophers have thought it impossible to deduce an 'Ought' from an 'Is,' or in general to deduce 'ethical sentences' from purely 'factual sentences.'¹ Some of these philosophers claim that this is due, not to any special feature of ethics, but to a general feature of logic, namely, its conservativeness:

A conclusion containing an 'ought' cannot (as a matter of logic) be derived from 'ought'-free premises. (The same, of course, goes for the other moral words.) Logic is conservative; the conclusions of a valid inference are contained within the premises. You don't get out what you haven't put in. Hence if 'ought' appears in the conclusion of an argument but not in the premises, the inference is not logically valid.²

This ban on inferences from Is to Ought on purely logical grounds — sometimes called 'the logical autonomy of ethics' — is widely accepted, but faces some well-known counter-examples. A.N. Prior, e.g., adduced

-
- 1 Charles Pigden suggests that the idea goes back as far as Cudworth. See C. Pigden, 'Logic and the Autonomy of Ethics,' *Australasian Journal of Philosophy* 67 (1989) 127-51. Pigden frames his discussion in terms of 'moral propositions' or 'moral judgements.' I find that terminology natural, but my main target, Toomas Karmo, frames his discussion in terms of 'ethical sentences,' so I shall do likewise.
 - 2 Charles Pigden, 'Naturalism,' in *Companion to Ethics*, P. Singer, ed. (London: Blackwell 1991), 423.

several apparently valid inferences from Is premises to Ought conclusions, including:

(ANP)

1. Tea-drinking is common in England.
Therefore,
2. Tea-drinking is common in England or all New Zealanders ought to be shot.³

More recently, Toomas Karmo has called attention to another sort of apparently valid inference from Is to Ought:

(TK)

1. Everything that Alfie says is true.
2. Alfie says that it ought to be the case that everyone is sincere.
Therefore,
3. It ought to be the case that everyone is sincere.⁴

Such counter-examples, however, do little to discomfit the believer in the logical autonomy of ethics (hereafter, 'autonomism'). About Prior's counterexample, Charles Pigden, exclaims, 'there is *something* wrong — something *odd* — about these inferences! You can't help feeling that Prior has taken himself (and his audience) in with a logical conjuring trick.'⁵ The trick, he explains, is that, although the conclusion contains an ethical disjunct, that disjunct is 'contingently vacuous.' We could substitute *any* expression of similar grammatical type, ethical or non-ethical, for 'All New Zealanders ought to be shot,' and the inference would still be valid. This suggests that the conclusion is not *really* an ethical conclusion at all.⁶ In view of this, the proper response to ANP is merely to formulate the thesis of logical autonomy more tightly as the thesis that there are no valid arguments from purely non-ethical premises to *non-vacuously* ethical conclusions.

3 A.N. Prior, 'The Autonomy of Ethics,' *Australasian Journal of Philosophy* 38 (1960), reprinted in *A.N. Prior: Papers in Logic and Ethics* A. Kenny and P. Geach, eds. (London: Duckworth 1976), 90-3

4 Toomas Karmo, 'Some Valid (but no Sound) Arguments Trivially Span the "Is"- "Ought" Gap,' *Mind* 97 (1988) 252-7, at 253. Karmo attributes this argument to Danny Mond. A similar argument is developed independently in Mark T. Nelson, 'Is it Always Fallacious to Derive Values from Facts?' *Argumentation* 9 (1995) 553-62.

5 Pigden, 'Logic and the Autonomy of Ethics,' 132

6 *Ibid.*, 133, ff. See also Pigden, 'Naturalism,' 424.

The same response does not work, however, against Karmo's counter-example: TK does not work by way of disjunction, e.g., and its conclusion cannot be faulted on grounds of vacuity. Karmo's response is to concede that the argument is valid, but argue that this sort of argument from Is to Ought can never be *sound*, and to formulate the thesis of logical autonomy still more tightly as the thesis that there are no *sound* arguments from non-ethical premises to non-vacuously ethical conclusions. Karmo's argument for this conclusion is ingenious and intriguing, but has recently been criticised by Stephen Maitzen.⁷ I think Maitzen is right, so I shall develop some different lines of criticism, calling into question some key elements of his argument, and raising a problem about the context in which autonomism is often deployed. I shall argue that, in that context, the most obvious responses to this problem are likely to be self-defeating. Before turning to this criticism, however, I must explain why autonomism might be thought to matter at all.

II Why the Logical Autonomy of Ethics Matters

Some philosophers used to think that the logical autonomy of ethics mattered because it bore on the debate over descriptivism versus prescriptivism in ethics. If ethical sentences were logically cut off from factual, descriptive sentences, this was one more reason for thinking that they themselves were not factual, descriptive sentences. That debate has gone off the boil, not least because the distinction between descriptive and prescriptive sentences has proven hard to draw, at least in any way that matters for ethics. If, for example, the distinction is drawn in terms of grammatical type, it faces the problem that some sentences in the indicative mood can function like prescriptions (e.g., 'You will do as I say'). If the distinction is drawn in terms of motivational or reason-giving potential, again it faces the problem that some apparently descriptive sentences can motivate or supply reasons for action (e.g., 'She is ill' or 'There is ice on the road ahead'). If the distinction is drawn in terms of the use of 'moral' and 'non-moral' words such as 'ought' and 'is,' the descriptive/prescriptive distinction just collapses into the is/ought distinction.

Autonomism still matters, however, because of its bearing on the debate over ethical skepticism, the claim that we can have no knowledge in ethics. In particular, autonomism matters because of the support it

7 Stephen Maitzen, 'Closing the "Is"- "Ought" Gap,' *Canadian Journal of Philosophy* 28 (1998) 349-66.

lends to 'justification skepticism' about ethics. The justification skeptic about ethics claims that we can't have ethical knowledge because we can't have adequate epistemic justification for our ethical beliefs. Ethical beliefs are truth-apt, and, for all anyone knows, some may even be true, but they want justification and fail to get it in the usual ways. How they fail is encapsulated in the following argument:

(ES)

1. Non-inferential knowledge of ethics is not possible.
 2. Sound arguments from factual premises to ethical conclusions are not possible.
 3. Good Inductive arguments from factual premises to ethical conclusions are not possible.
 4. Ethical knowledge is possible only if: non-inferential knowledge of ethics is possible, or sound arguments from factual premises to ethical conclusions are possible, or good inductive arguments from factual premises to ethical conclusions are possible.
- Therefore,
5. Ethical knowledge is not possible.

The conclusion of this argument is skepticism about ethics in particular, but the argument itself follows a general skeptical strategy.⁸ That strategy requires us to distinguish sharply between two domains, and to identify one as epistemologically *unproblematic*, and the other as epistemologically *problematic*. Different accounts may be given of what makes a domain 'unproblematic,' but the usual ones are that we have empirical access to that domain, that we are related to the alleged facts in that domain in the right causal or explanatory way, and that there is nothing essentially mysterious about our knowledge of that domain. A domain is 'problematic,' on the other hand, when it is hard to see how we have direct empirical access to it, when we are 'cut off' causally or explanatorily from its alleged facts, or when knowledge of it would be

8 Berkeley gives this argument for his skepticism about physical objects, though the form of his argument is not so explicit. See *A Treatise Concerning the Principles of Human Knowledge*, sections 18-21. An explicit version of the argument is applied to mind-independent physical objects in W.T. Stace, 'The Refutation of Realism,' *Mind* 53 (1934) 145-55. For a discussion of the argument as applied to physical objects, the past and personal identity, see A.J. Ayer, *The Problem of Knowledge* (Harmondsworth: Penguin 1956), 75-81. Something like this argument is applied to ethics in Walter Sinnott-Armstrong, 'Moral Skepticism and Justification' in *Moral Knowledge? New Readings in Moral Epistemology*, W. Sinnott-Armstrong and M. Timmons, eds. (Oxford: Oxford University Press 1996), 3-48, especially 10-12.

essentially mysterious. The skeptic then requires us to vindicate claims to knowledge about the problematic domain in terms of the unproblematic domain, where vindication is usually thought to require definitional reduction to, or logical inference from, materials in the unproblematic domain. Here, the skeptical argument distinguishes between a purely 'non-ethical' domain, and a purely 'ethical' domain.⁹ It holds that the non-ethical domain is natural, factual and epistemologically unproblematic, but that the ethical domain is problematic.

The problematic character of the ethical domain is affirmed immediately in premise (ES1), which denies that directly justified knowledge of ethics is possible. ('How could one have direct knowledge of anything other than natural, everyday facts? And how could one have non-inferential knowledge of the ethical, when all parties to the debate admit that the ethical supervenes on the non-ethical?') This implies that, if we are to have any ethical knowledge, it must somehow piggyback on other sorts of knowledge. Premises (ES2) and (ES3) deny that ethical knowledge *can* piggyback on other sorts of knowledge, because ethical sentences cannot be inferred, deductively or inductively, from epistemologically unproblematic sentences about the domain of the non-ethical. Premise (ES4) might be thought unnecessary, but it is included to make the argument explicitly logically valid, and can be read as saying no more than 'These are the only ways that ethical beliefs could be epistemically justified.' The relevance of autonomism to this skeptical argument should be clear: the reason usually advanced for premise (ES2) just is the logical autonomy of ethics.¹⁰ There can't be sound arguments from factual premises to genuinely ethical conclusions, if there can't be valid, non-vacuous inferences from factual premises to ethical conclusions. This alone is reason to care about whether alleged counter-examples to autonomism, such as TK, really work.¹¹

9 I doubt that these domains can be non-arbitrarily partitioned, but I shall ignore this doubt for the sake of argument.

10 Pigden explicitly recognizes this connection between logical autonomy and skeptical arguments. See Pigden, 'Logic and the Autonomy of Ethics,' 130, n. 6.

11 As an anonymous referee for this journal pointed out to me, it should be noted that ethical autonomism is neither a necessary nor a sufficient condition for ethical skepticism. One might accept autonomism but hold that non-inferential knowledge of ethics is possible. Equally, one might reject autonomism but hold that ethical knowledge fails for some other reason. For a statement of the latter position, see Walter Sinnott-Armstrong, 'From 'Ought' to 'Is' in Moral Epistemology,' *Argumentation* 14 (2000) 159-74.

III Karmo's Argument against Counter-example TK

Karmo argues that counter-examples such as the following *don't* work:

(TK)

1. Everything that Alfie says is true.
 2. Alfie says that it ought to be the case that everyone is sincere.
- Therefore,
3. It ought to be the case that everyone is sincere.

On the face of it, this argument has Is-premises and an Ought-conclusion and is logically valid; if its premises be true, its conclusion cannot be false. It cannot be sound, however, because as soon as its conclusion is true, at least one of its premises *becomes* ethical.¹² Karmo's argument for this puzzling-sounding claim depends on two assumptions and three definitions. The first assumption is that whether a sentence is an ethical sentence is a purely contingent matter; on this view, the same premise can be ethical in some possible worlds, but non-ethical in others.¹³ The definitions are: 1) a sentence S is ethical in a possible world w just in case S is true in w with respect to one ethical standard, and false in w with respect to another ethical standard; 2) an ethical standard is a complete assignment of truth values to all uncontroversially ethical sentences; and 3) a sentence is uncontroversially ethical just in case all parties to the logical-autonomy-of-ethics debate would unite in calling it ethical.¹⁴ The final assumption is that:

just as some one possible world is the actual world, so some one ethical standard is the correct ethical standard. When people simply say, 'Sentence S is true,' we take them to mean 'S is true in the actual world with respect to the correct ethical standard.' When people simply say, 'S is true in world w,' we take them to mean 'S is true in w with respect to the correct ethical standard.'¹⁵

With these assumptions and definitions in place, Karmo formulates his claim thus:

12 This is misleading; it is not exactly a matter of *becoming* ethical. As I shall explain, the change is not over time, but over possible worlds.

13 Here he follows Lloyd Humberstone, 'First Steps in Philosophical Taxonomy,' *Canadian Journal of Philosophy* 12 (1982) 467-78.

14 Karmo, 'Some Valid,' 254

15 Ibid.

In general, if sentences S_1, \dots, S_n (where $n > 0$) entail sentence $S(n + 1)$, then for any possible world w in which $S(n + 1)$ is ethical, if all of S_1, \dots, S_n are true in w , then at least one of S_1, \dots, S_n is ethical in w . (Proof: since all of S_1, \dots, S_n are true in w — that is, are true in w with respect to the correct ethical standard E — $S(n + 1)$ is true in w with respect to E . If $S(n + 1)$ is to be ethical in w , then there must be some alternative ethical standard E' with respect to which $S(n + 1)$ is false in w . Since replacing E with E' changes the truth value of $S(n + 1)$ in w , and $S(n + 1)$ is entailed by S_1, \dots, S_n , replacing E with E' changes the truth value in w of at least one of S_1, \dots, S_n ; so at least one of S_1, \dots, S_n is ethical in w .)¹⁶

We can illustrate his approach in less abstract terms by showing how it applies to TK. Let us assume, for the sake of argument that:

- i) TK is valid;
- ii) (TK2) is non-ethical (as it appears to be);
- iii) (TK3) is ethical.

To determine whether TK is a sound Is-to-Ought argument, and hence a counter-example to the logical autonomy of ethics, we must ask:

- iv) Is (TK1) true?
- v) Is (TK2) true?
- vi) Is (TK1) non-ethical?

And we must get 'yes' answers to all three. Let us consider several of Karmo's own scenarios, and the results they yield, according to him.

Scenario 1: in world w_1 , Alfie says only one thing: 'Tar is inflammable.' In that case:

- iv) Is (TK1) true? Yes, because the only thing that Alfie says is true.
- vi) Is (TK1) non-ethical? Yes, because it is true regardless of what ethical standard obtains. (Kantians, utilitarians, and egoists, e.g., can all agree in w_1 that (TK1) is true.)
- v) Is (TK2) true? No, because it represents Alfie as saying something he didn't say.

Result: in world w_1 , TK is a valid argument from Is to Ought, but is not sound, because Alfie doesn't make any ethical pronouncements. Let us consider, therefore, a world in which he does.

16 Ibid, 256.

Scenario 2: in world w_2 , Alfie says only one thing: 'It ought to be the case that philosophers are vegetarians.' In that case:

- iv) Is (TK1) true? Here, it depends both on what he says and on what the correct ethical standard is. If the correct ethical standard is Schweizerian, the answer is yes; if it is, say, Schwarzeneggerian, the answer is no.
- vi) Is (TK1) non-ethical? No, because, as above, (TK1) will be true according to some ethical standards and false according to others. In w_2 , (TK1) is an ethical sentence.
- v) Is (TK2) true? No, because it represents Alfie as saying something he didn't say.

Result: in world w_2 , TK is a valid argument, but it has at least one Ought-premise; moreover, it is unsound. Suppose we try to fix the unsoundness by creating a scenario in which Alfie both utters an ethical sentence and *does* say what the argument represents him as saying?

Scenario 3: in world w_3 , Alfie says only one thing: 'It ought to be the case that all philosophers are sincere.' In that case:

- v) Is (TK2) true? Yes, because it represents Alfie as saying what he did in fact say.
- iv) Is (TK1) true? Again, it depends on what the correct ethical standard is. If the correct ethical standard is Kantian, the answer is yes; if Machiavellian, the answer is no.¹⁷
- vi) Is (TK1) non-ethical? No, because, as above, (TK1) will be true according to some ethical standards and false according to others. In w_3 , (TK1) is an ethical sentence.

Result: in world w_3 , TK is a valid argument (and maybe even sound), but it has at least one Ought-premise, so it is not a sound Is-to-Ought argument.

According to Karmo, such scenarios illustrate the upshot of his overall argument: there are trivially valid arguments from Is to Ought, but there are no sound arguments from Is to Ought, because, given Karmo's theoretical machinery, the 'ethicality' of any sound ethical conclusion works backward, as it were, through the logical structure of a deductive argument and necessarily infects at least one of the premises. Thus, Karmo supplies a tidy proof of (ES2), 'Sound arguments from factual premises to ethical conclusions are not possible,' a key premise in ES, the

17 I am assuming here simplistic stereotypes of Kant and Machiavelli; nothing hangs on this.

general argument for justification skepticism in ethics. It is not clear whether Karmo intends his argument for autonomism also to support ethical skepticism; he says nothing about it. It is clear, however, that autonomism lends itself readily to such a use. Since the time of Hume, ethical autonomism has commonly been perceived as raising problems for ethical knowledge, and Lloyd Humberstone, from whom Karmo borrows his theoretical machinery, explicitly raises the epistemological implications of such 'inference barriers':

Suppose there are two boxes in front of me, box A and box B, either or both of which may be filled in any of various (independent) ways. But box B is locked, while box A can be readily opened. Problem: to divide statements about what's in the boxes into two classes, the A-ish and the B-ish in such a way as to render true the intuitively plausible claim that no amount of inspection of the contents of box A will give you any information about box B.¹⁸

The epistemologically unproblematic domain of fact is like the open box A; the epistemologically problematic domain of ethics is like the locked box B. If no sound inferences exist from premises about the contents of the fact box to conclusions about the contents of the ethics box, the contents of the ethics box seem forever beyond our ken.

IV Evaluation of Karmo's Argument

The first thing to notice about Karmo's treatment of counter-example TK is that it does not simply show, as it happens, that no sound Is-to-Ought arguments exist. It shows that they *could not exist*, by legislating against their very possibility. If we are told, as a consequence of some definitions, that any set of true sentences logically entailing an ethical sentence *must* contain at least one ethical sentence, we are not being given a *reason* for Hume's ban; we are being given Hume's ban under another guise. Arguably, this begs the very question being debated, because it tests sentences for 'ethicality' in terms of their entailments *jointly* with other sentences. A fairer, non-question-begging approach would test sentences individually. It would then remain an open question until after those sentences were assembled, whether they constituted a sound Is-to-Ought argument. At the very least, this implies that all the weight of the debate must now be borne by whatever reasons may exist for preferring Karmo's definitions and assumptions in the first place.

18 Humberstone, 'First Steps,' 477. Note, however, that, instead of wanting his tools to *support* ethical skepticism, Humberstone actually wants them to *presuppose* it!

So, what about Karmo's definition of sentences as ethical just in virtue of differential entailments with respect to alternative ethical standards? I suggest that it has several problems. First it is inconsistent (or at least in tension with) Karmo's other definitions. Karmo's definition (1), of 'ethical sentence,' offers a *truth-value dependent* way of determining what is ethical. Karmo's definitions (2) and (3), however, define 'ethical standards' in terms of 'uncontroversially ethical sentences,' and 'uncontroversially ethical sentences,' in turn, in terms of something *truth-value independent*, namely the disputants' uniting in calling them ethical. The tension can be brought out by asking whether the disputants can identify these sentences as ethical just by looking at them, without checking entailments and truth-values. If they can, this suggests that they have some truth-value independent criterion for ethical sentences per se, such as 'having ethical content.' But if they cannot, how is their partitioning of sentences into ethical and non-ethical supposed to get started?¹⁹

Second, this definition is odd, precisely because its criterion for being an ethical sentence is truth-value dependent, and thus bears little relation to the content of the sentences in question. The oddness of this can be brought out by noting that Karmo's criterion allows him to be completely indifferent about *which* of the premises of argument TK is the ethical sentence; it insists only that one of them must be. For the purposes of argument, we assumed that (TK2) was non-ethical, so Karmo's theoretical machinery automatically picked (TK1) as ethical. If, on the other hand, we had assumed (TK1) was non-ethical, it would have picked (TK2) as ethical, just as automatically. Karmo can reply that this is part and parcel of his assumption that sentences are only contingently ethical, but the odd implications of that previously bland-seeming assumption are now plainly visible, and this apparent indifference to content may lead us to suspect that we have now lost our grip on the ethical/non-ethical distinction altogether.

Third, Stephen Maitzen has recently shown that it cannot be correct, because it yields wrong answers about sentences such as:

B2. Either no ethical sentence is true, or torturing babies just for fun is wrong.²⁰

19 Thanks to an anonymous referee for helping me to see this.

20 Maitzen frames his discussion in terms of 'moral propositions' instead of 'ethical sentences,' but I have changed it to 'ethical sentences,' to fit in with Karmo's wording.

Maitzen argues that any sentence that is logically compatible with moral nihilism, the thesis that no ethical sentence is true, cannot be an ethical sentence.²¹ Sentence (B2), however, is logically compatible with nihilism. J.L. Mackie, e.g., could easily accept it: 'Yes, at least one of its disjuncts is true. (The first one!)' Karmo, on the other hand, would be forced by his theoretical machinery to classify (B2) as an ethical sentence, since it comes out true on some ethical standards and false on others. All non-nihilistic ethical standards have to hold the first disjunct as false, so for them, the truth or falsity of (B2) overall will hang on the truth or falsity of the second disjunct. Ethical standards condemning baby torture will hold (B2) true; ethical standards condoning it will hold it false. This is all that it takes for (B2) to come out as an ethical sentence according to Karmo and to point up the inadequacy of his definition.²² I think that Maitzen's critique here is correct, but even if it weren't, Karmo's approach would face other problems. In particular, it has an unwelcome consequence in the context of ES, the argument for ethical skepticism.

Consider the following sort of argument, which might be given by a conventional moralist who is also an admirer of Stalin:

(JS)

1. If Stalin authorized the Katyn Massacre then Stalin was evil.²³
2. It is not the case that Stalin was evil.

21 A similar criterion is developed in Nelson, 'Is it Always Fallacious to Derive Values from Facts?'

22 Using this result, Maitzen also manages to refute Karmo's overall conclusion by constructing a counterexample: a simple Is-to-Ought argument which can plausibly be regarded as sound, but whose premises can be *shown* to be purely non-ethical:

B1. Some ethical sentences are true.

B2. Either no ethical sentence is true, or torturing babies just for fun is wrong.
Therefore,

B3. Torturing babies just for fun is wrong.

This argument is valid via disjunctive elimination; the conclusion is an 'uncontroversially ethical sentence'; (B2) has already been shown to be an ethical sentence, and is true according to both nihilistic and conventional standards. Maitzen argues that (B1) also is a non-ethical sentence, on the grounds that we must distinguish between 'sentences of ethics' and 'sentences about ethics.' Sentences of ethics entail that some particular moral property is possessed by some object(s) or other. (B1) entails no such thing, so it is better classified as a sentence about ethics, on a par with, say, "'Morality'" often goes by the label "'ethics.'" See Maitzen, 'Closing the "Is"- "Ought" Gap,' 361.

23 The Katyn Massacre was the execution by the NKVD in April-May, 1940 of approximately 25,700 members of the Polish intelligentsia, including 14,700 prisoners of war, in a forest near the village of Katyn, near Smolensk, Russia.

Therefore

3. It is not the case that Stalin authorized the Katyn Massacre.

Among other things, this shows how an argument with apparently ethical premises can validly entail an apparently non-ethical conclusion. What if we turn it inside out, as follows?

(JS*)

1. Stalin authorized the Katyn Massacre.

Therefore

2. It is not the case that both *If Stalin authorized the Katyn Massacre, then Stalin was evil* and *It is not the case that Stalin was evil*.

This argument also is valid.²⁴ Moreover, its sole premise is apparently true, or at least verifiable in some of the ways that historical claims typically are (e.g., via a signed piece of paper or the testimony of an eye witness). To determine, however, whether JS* is a sound Is-to-Ought argument (hence a counter-example to the logical autonomy of ethics), we must ask:

- vii) Is (JS*2) ethical?
- viii) Is (JS*1) non-ethical?

To answer (vii), by Karmo's lights we must ask whether (JS*2) comes out true on at least one ethical standard and false on another. A tough-minded act-utilitarian, for example, can deny that *If Stalin authorized the Katyn Massacre, then Stalin was evil*, so there is at least one ethical standard according to which (JS*2) is true. On the other hand, we can conceive of an ethical standard according to which both *Stalin was not evil*, and *anybody who authorized the Katyn massacre would be evil* (perhaps a standard combining belief in non-violence with belief in the virtue of Georgians). Therefore, there is at least one ethical standard according to which (JS*2) is false. Therefore, according to Karmo's definition, (JS*2) is indeed an ethical sentence, and we appear to have a valid argument to an ethical conclusion.²⁵

Regarding (viii), however, won't the autonomist just run the same argument that Karmo ran for TK? Won't she argue that, since (JS*1)

24 Because p entails $\sim((p \rightarrow q) \ \& \ \sim q)$.

25 Moreover, the separate conjuncts in (JS*2) are, on Karmo's criterion, clearly ethical sentences in w , so it would be very odd if their conjunction ceased to be an ethical sentence in w ! Any argument that (JS*2) is not an ethical sentence must face up to this odd implication.

entails (JS*2), and the truth value of (JS*2) varies relative to different ethical standards, then the truth value of (JS*1) also varies relative to different ethical standards, hence (JS*1) is also an ethical sentence? It may *look like* a non-ethical sentence, but won't the autonomist just shrug her shoulders and say that the lesson we learned from counter-example TK was that such appearances can be deceiving?

The autonomist *could* make that response, but it would be undesirable in the context of one of the main arguments in which autonomism importantly figures, namely ES, the argument for justification skepticism in ethics. One of the key assumptions of that argument, as noted above, is that the domain of the ethical is epistemologically problematic. If (JS*1) is an ethical sentence after all, and it is also the sort of claim that can be verified by whatever means historical claims are normally verified (appeal to signed documents, eyewitness reports, forensic science), then we have an example of an epistemologically unproblematic ethical sentence. Far from aiding the argument for justification skepticism in ethics, autonomism à la Karmo undermines one of its key assumptions. Moreover, this criticism generalises: even if JS* is not a serious threat to autonomism or skepticism, Karmo's strategy, when applied to other alleged Is-to-Ought arguments, would risk creating similar problems for the skeptical deployment of autonomism.²⁶

The autonomist, of course, has a reply here: 'Yes, on my view, (JS*1) does count as an epistemologically unproblematic ethical sentence, but before the anti-skeptic begins to gloat, he should realise that (JS*1) is not the sort of belief that would improve anyone's ethical knowledge very much. There are ethical sentences, and then there are ethical sentences: the sort that you have adduced in (JS*1) are epistemologically unproblematic; the epistemologically problematic sort are the ones that offer guidance, evaluate people or actions, or rely on certain distinctively ethical words or concepts. (JS*1) is clearly of the first sort, and doesn't improve our knowledge of the second sort at all — and these are the ones that matter!'

The autonomist can make this reply, but it is not without philosophical cost: it is tantamount to an admission that the latter are the *real* ethical sentences after all, hence that there are some further criteria for being an ethical sentence. This, in turn, casts doubt on Karmo's definition of

26 For example: Searle's Is-to-Ought argument, which proceeds from one epistemologically unproblematic non-ethical sentence and a number of *ceteris paribus* qualifiers. See 'How to Derive "Ought" from "Is,"' *Philosophical Review* 63 (1964) reprinted in *The Is/Ought Question*, W.D. Hudson, ed. (London: Macmillan 1969), 120-34.

'ethical sentence' purely in terms of logical relations to ethical standards. If the autonomist makes this reply, she undermines her rejection of counter-example TK and leaves autonomism open to refutation. If she doesn't, she undermines one of the key assumptions of the larger skeptical argument, namely, the claim that the domain of ethics is essentially epistemologically problematic. Either way, it's more bad news for the logical autonomy of ethics, at least in its skeptical applications.²⁷

Received: June 2005

Revised: March 2006

27 The writing of this essay was supported by the Institute of Faith and Learning at Baylor University. I thank the Institute and the Baylor Philosophy Department for their support, and Jennifer Jackson, Robin LePoidevin, Andrew McGonigal, Stephen Maitzen, Scott Shalkowski, and two anonymous referees for this journal for helpful comments and discussion.