

GIVING UP HUME'S GUILLOTINE

Aaron Wolf

The appealing principle that you can't get an 'ought' from an 'is', sometimes called Hume's Guillotine (or, more soberly, Hume's Law), faces a well-known challenge: it must give a clear account of the distinction between normative and descriptive sentences while dodging counter-examples. I argue in this paper that recent efforts to answer this challenge fail because the distinction between normative and descriptive sentences cannot be described well enough to be of any help. As a result, no version of the principle is both true and adequate. Yet a different principle—that no normative terms are synonymous with any descriptive terms—can satisfy much of the motivation we had for defending Hume's Guillotine in the first place. I show briefly how this second principle can explain why is-ought arguments are invalid, block some attempted deductive justifications of beliefs with normative terms, and create an explanatory advantage for moral non-cognitivism. Just so, giving up Hume's Guillotine is the prudent thing to do.

Keywords: no ought from is, logic, Hume, A.N. Prior, G.E. Moore

1. Introduction

There's a difference between truth and relevance. Hume's Guillotine—the idea that you can't get an 'ought' from an 'is'—rests on a distinction between normative and descriptive sentences.¹ This is because Hume's Guillotine is an inference barrier, and inferences deal in sentences. So, to say that you can't get an 'ought' from an 'is' is really to say that you can't validly infer any normative sentences from any descriptive ones. Though it was long out of favour, the Guillotine has enjoyed a resurgence among technically-minded philosophers who hope to prove that some version of it is true and thereby to silence the doubters. Proving something in a formal language is not enough to vindicate the principle, though. In order to be a *bona fide* Hume's Guillotine, a principle must be a guide for ethics as we actually do it in natural languages. That is, it needs to govern real-life normative assertion. I suggest this as a criterion of success on any attempt to defend the view. Without this kind of relevance, there is little reason to care about Hume's Guillotine.

Because of powerful counter-examples from A.N. Prior, anyone defending the Guillotine must give an account of which sentences are normative and which are descriptive. Getting this right is crucial. It's one thing to prove a result in a language, and it's quite another to prove a result in one that can serve as an appropriate meta-language for the ethics we actually do.

¹Hume's Guillotine is known by many other names. I borrow this more vivid label from Max Black [1964].

I argue in this paper that the prospects for proving an acceptable version of Hume's Guillotine are dim, because the manoeuvring required to secure a true principle alienates it from real-life normative assertion. This is a reason to think that the effort to defend the Guillotine with formal methods is failing. That project may ultimately succeed, but too little attention goes to creating an adequate model of ethics. I hope to show the cost of that neglect.

In what follows I'll show that each successful strategy for beating Prior's counter-examples involves a faulty distinction between normative sentences and descriptive ones. Irrelevance across the wide range of views strongly suggests that the problem lies with this sentence-level distinction. Some perfectly natural divisions may not be clean enough for philosophical use, so choosing the right distinction on which to stake a claim makes all the difference.

Many have argued that the Guillotine is problematic, and have said no more. I want to make the additional point that much of what's desirable in it can be gained with a principle about words rather than about sentences: no descriptive terms are synonymous with any normative terms. Call this view *Moore's Law* because of its connection with the naturalistic fallacy. The problems for Hume's Guillotine arise because it's an inference barrier and therefore requires a clear distinction between sentences. Moore's Law separates the normative from the descriptive without needing that distinction. The inference barrier is replaced by a semantic one, and a term-level distinction takes the place of the sentence-level one. Insofar as it's worth having a methodological constraint like the Guillotine, Moore's Law offers a nice replacement.

Given how often the Guillotine is invoked, it's clear that many do want such a rule for ethics. But if what I argue is right, those of us who treat the Guillotine as a philosophically secure way to correct students and to rebuke opponents will have to think again. Those who have hoped to get even more mileage out of it by enlisting it to show some further result will have to re-work their arguments.² The good news for all of us is that Moore's Law can deliver these goods as well.

Relevance means that a guiding principle like the Guillotine should be a rule for a game that people are really playing. An inference barrier couched in a language we're not speaking is no more significant than a barrier to inferences we're not making, and a false principle governs nothing. Since the versions of the Guillotine on offer require faulty ideas about what counts as normative, they're rules for the wrong game.

This concern with relevance is not the only one out there, however. Gerhard Schurz and Charles Pigden have a dispute over whether or not Hume's Guillotine is relevant in the sense of having other metaethical implications, particularly whether or not it implies anything about non-cognitivism or naturalism.³ This is a different question. Someone like Pigden who believes that it's metaethically neutral could also believe consistently that it's a guiding rule for ethics and ethical inferences. Pigden suggests he does think this,

²Sinnott-Armstrong [2000] and Huemer [2005] invoke the Guillotine to defend views about moral epistemology, for example.

³See Pigden [2010b: 217–38] and Schurz [2010b: 239–42].

when he says that people trying to infer normative conclusions from descriptive premises are making a significant mistake [2010a: 26]. So metaethical relevance is not the same thing as first-order ethical relevance.

Although I doubt that the sentence-level distinction can be made well enough to deliver a relevant version of the Guillotine, I think there is a real difference and that some sentences are clearly normative or descriptive. That's compatible with thinking that explaining the difference is hard to do, and it's also consistent with allowing difficult borderline cases. Indeed, the basic problem that motivates the views we'll see below is brought out by such a case.

In the next section, we'll see the dilemma for Hume's Guillotine that forces its defenders to give specific accounts of normative sentences. In sections 3 through 5, I argue that efforts to frame an inference barrier around these accounts fail to produce a suitable version of the view. We conclude in section 6 by observing how Moore's Law can stand in for the Guillotine, and offer at least one improvement on it. Just so, Moore's Law is well-positioned to be the guiding principle we wanted.

2. Prior's Dilemma

Prior [1960] thought he could show that Hume's Guillotine was false in a way that most other attempts could not. A popular strategy for arguing against the view went like this: identify a normative sentence, then argue that it follows from some set of descriptive sentences. This was John Searle's [1964] approach in his example where Jones promises to pay Smith five dollars.⁴ From the descriptive premise that Jones uttered the words of the promise, we can supposedly infer the normative conclusion that Jones must pay Smith five dollars. A cost of this strategy is that its success depends on whether the sentences really are normative and descriptive as alleged, and critics of Searle were quick to point this out. Prior's argument is especially powerful because it doesn't require consensus on which its crucial sentence is.

Start by considering this basic statement of the Guillotine.

Naïve HG: There are no satisfiable valid arguments from descriptive premises to normative conclusions.⁵

The argument against Naïve HG then goes as follows [Prior 1960]. Take the valid argument, A:

- (A-i) Tea-drinking is common in England.
- ∴ (A-ii) Either tea-drinking is common in England or all New Zealanders ought to be shot.

⁴Other instances of the strategy can be found in Black [1964], Castañeda [1973], and Geach [1977].

⁵Being satisfiable means the premises are consistent. We need this restriction because anything follows from a contradiction.

The premise of A is descriptive. So if the conclusion is normative, then Naïve HG is false. But since it need not be true that all New Zealanders ought to be shot if the conclusion is to be true, one might suspect that it's not normative. Suppose then that it's descriptive, and consider a second valid argument, B:

- (B-i) It's not the case that tea-drinking is common in England.
- (B-ii) Either tea-drinking is common in England or all New Zealanders ought to be shot.
- ∴ (B-iii) All New Zealanders ought to be shot.

Again, the premise is clearly descriptive. And, plausibly, the tea-drinking disjunction must be either normative or descriptive. If we say it's normative, then argument A is a counter-example to Naïve HG. But if we say it's descriptive, then B is a counter-example. Hence the dilemma.

The challenge for defenders of Hume's Guillotine is to get around Prior's dilemma. By my count, there are three ways of doing this.

- (1) Accept that A is a valid is-ought inference, but replace Naïve HG with something that doesn't apply to A (the term approach).
- (2) Call the tea-drinking disjunction normative as a premise but descriptive as a conclusion (the relativity approach).
- (3) Say that the tea-drinking sentence is neither normative nor descriptive (the fragility approach).

I'll discuss these strategies in order below.

3. First Approach: Terms

Normative sentences might be distinguished from descriptive ones by their terms. Most words are clearly one or the other. On this strategy, a sentence is normative just in case it contains a normative term, like 'ought', 'should', 'may', 'good', 'right', and 'wrong', and descriptive just in case it doesn't. Call this the *term approach* to normative sentences.

The term approach handles the B-argument of Prior's dilemma where the disjunction is a premise, because it has an 'ought'. But it doesn't block the A-argument where the disjunction is the conclusion, since its premise has no normative terms. Yet Frank Jackson [1974] observes that there is something odd about the A-argument: the 'ought to be shot' in the conclusion is replaceable by any grammatical expression whatsoever—'are friendly' or 'prefer rugby to curling'—and the argument remains valid, since it's an instance of the rule of addition. Call a term in an argument *vacuous* just in case each occurrence is uniformly replaceable without prejudice to the

inference's validity.⁶ This insight gives way to a version of Hume's Guillotine that avoids the looming counter-example.

Jackson's principle rules out, from the scope of the view, any argument with vacuous terms [Jackson 1974: 93].

JHG: No argument from premises lacking normative terms to a conclusion with normative terms is valid, unless one of the argument's terms is vacuous.

Prior's A-argument is no counter-example to this version of Hume's Guillotine, since the 'ought' of the tea-drinking disjunction is vacuous. But the restriction against any vacuous term goes too far. Jackson's principle places outside of its scope any argument that is valid in virtue of logical form. Think of *modus ponens*, for example: in an inference like *if ϕ -ing would maximize utility then it's the right thing to do, ϕ -ing would maximize utility; therefore it's the right thing to do*, the terms of both the antecedent and consequent are vacuous, so it's outside the scope of the view. Thus the only arguments to which JHG properly applies are those that are valid just in virtue of the meaning of their terms.⁷

It's not difficult to see why this goes too far. Philosophical ethics is in the business of giving formally valid arguments like the one above, as well as informal ones. So if JHG has nothing to say about arguments that follow *modus ponens*, disjunctive syllogism, etc., then it cannot be a guiding principle of ethics as we actually do it.

Pigden [1989] extends Jackson's strategy by framing a version of the Guillotine around vacuity, but only when it occurs in the conclusion. Roughly, Pigden's says this:

PHG: No argument from premises lacking normative terms to a conclusion with normative terms is valid, unless the normative terms in the conclusion are vacuous.

One immediate concern about PHG (which also affects JHG) is that it's under-specified—it fails to distinguish use from mention. A sentence might contain a term in either way, so the term approach really covers two distinct ideas. There is a *narrow* claim which says that a sentence is normative just in case it uses a normative term, and a *wide* one which says that a sentence is normative just in case it uses or mentions a normative term. We'll tackle the narrow version first:

NHG: No argument from premises not using normative terms to a conclusion that uses normative terms is valid, unless the normative terms in the conclusion are vacuous.

This appears to be what Pigden had in mind. He does not explicitly consider the use/mention distinction. But he speaks of predicates occurring, not of

⁶Cf. Pigden [1989: 134].

⁷Jackson appears to agree. See his [1974: 91].

individual constants. So it's safe to assume that this is the principle he would choose [1989: 136].

The trouble with NHG is that while it avoids Prior's dilemma, it's susceptible to another kind of counter-example. To see how the example works, consider this wholly descriptive argument:

(C-i) 'Clark Kent' names the reporter from Krypton who writes for the Daily Planet.

∴ (C-ii) Clark Kent writes for The Daily Planet.

Argument C is truth-functionally valid, in the sense that the truth of the premise guarantees the truth of the conclusion. It's not formally valid, but restricting the Guillotine to just formally valid arguments won't do, as we've seen. And although the conclusion may not be derivable by syntactic inference rules, it does clearly follow from the naming convention established by the premise. Once it's established, we are licensed to use the name. Furthermore, 'Clark Kent' is clearly non-vacuous. Replacing the conclusion with 'Lois Lane writes for the The Daily Planet' would invalidate the argument. Now consider argument D, which trades on the same kind of inference:

(D-i) 'Ought to' names the relation that holds between Lois and the action-type *donate to charity*.

∴ (D-ii) Lois ought to donate to charity.

If you're thinking that the premise of this argument is normative, I'm inclined to agree. It's just a roundabout way of saying that Lois ought to donate to charity. In the absence of any particular view about what makes for a normative sentence, the D-argument is not an especially strong candidate for an is-ought derivation.⁸

But responding to Prior's dilemma does require a particular view of normative sentences, because Naïve HG looks false. For NHG, it's the narrow term approach. It says that (D-i) is descriptive, (D-ii) is normative, and the 'ought' in the conclusion is non-vacuous. So D is a counter-example to NHG. And, as C shows, it's not a contrived one-off argument, either. It's not hard to imagine genuine instances of this kind of reasoning. Perhaps an enemy of Superman has obtained a government list of known superhero identities, hoping to threaten those about whom Superman cares. Or perhaps a powerful computer has generated a list of possible actions for Lois that carry extremely high expected utility; but due to limitations in the programming language, results are given metalinguistically. Given the information in the premise, the inference to the conclusion is an easy one, but that doesn't make D any less of a counter-example.

Defenders of Pigden will be quick to point out that he has given a proof, so there must be something amiss with the counter-example. But the trouble really lies with his proof. The problem is that it can't show NHG, because it fails to distinguish use from mention. Pigden hopes to prove that a predicate

⁸I imagine this is why no one else I know of has tried this sort of counter-example.

'cannot occur non-vacuously in the conclusion of a valid inference unless it appears among the premises' [1989: 136]. As we've seen, predicates can occur in both ways. Argument D reveals that the proof shows at most that a term cannot be used in the conclusion of a valid argument unless it is used or mentioned. But that's not the same as demonstrating NHG; that's showing something about the wide term view, which is weaker—the disjunction is crucial.

Objectors might try another response. In a recent commentary on a different class of counter-examples—which infer normative conclusions from descriptive premises by claiming that some group of sentences is true and by including some normative-looking sentence among them, e.g. *everything Clark believes is true, and Clark believes Lois ought to donate to charity; therefore Lois ought to donate to charity*—Pigden objects that they 'can be accused of concealing the moral content of one of their premises behind the truth predicate' [2010a: 36]. Likewise, I can be accused of concealing the moral content of the premise in the D-argument behind the naming convention it establishes, and I readily admit the inference does that. The question is whether this bit of linguistic trickery makes a difference to NHG. In his commentary, Pigden speaks of a distinction between 'formally' moral and 'substantively' moral sentences [2010a: 34–8]. Some may therefore object that even if the premise of D is formally descriptive by the rules of NHG, it is nonetheless substantially normative and so it fails to undermine the spirit of the Guillotine.

Yet this sort of objection gets things backwards. Advocates of the term strategy for combating Prior's dilemma are not entitled to fall back upon some more general or 'substantial' version of the view if the formal candidate fails. Since the predicate approach allows Prior's tea-drinking disjunction as normative, Naïve HG is false by its lights. It's this trouble with the naïve claim that forces defenders of the Guillotine to come up with something more sophisticated. If one of the sophisticated principles succeeds, that just is the substantial view. If it fails, it's time to go back to the drawing board, since there is nothing defensible to fall back on. Those who want to defend Hume's Guillotine cannot have it both ways. Either the replacement for Naïve HG is the genuine article, or it isn't.

The D-argument also undermines Schurz's [2010a] defence of the Guillotine, which categorizes sentences as either purely normative, purely descriptive, or mixed. Because he wants to prove versions of it for a range of deontic logics, Schurz identifies sentences by deontic operators rather than by logical constants. But the basic approach is the same as Pigden and Jackson's: the words themselves make the difference. A sentence is purely normative just when the entire thing lies inside the scope of a normative operator, purely descriptive just when there is no use of a normative operator, and mixed just when the operator binds only part of the sentence [Schurz 2010a: 200–1]. So a sentence like 'Lois ought to donate to charity' will be purely normative, and a sentence like 'snow is white' will be purely descriptive. Because only its right-hand disjunct is in the scope of a normative operator, Prior's tea-drinking disjunction comes out mixed.

Given these categories, there are three possibilities for a valid is-ought inference. We might have (a) purely descriptive to mixed (like Prior's

A-argument where the disjunction is the conclusion), (b) mixed to purely normative (like Prior's B-argument where the disjunction is a premise), or (c) purely descriptive premises to purely normative. Schurz appears to think that others have shown that no non-trivial type-(c) arguments are valid, so he leaves out this possibility and he focuses on blocking types (a) and (b).⁹

But ignoring (c) is a mistake—the D-argument is a valid instance of it by Schurz's lights. The naming convention for Lois given by the premise must be interpreted here as purely descriptive. Although there is a mention of 'ought' and although operators are defined syntactically, it can hardly be thought of as an active normative operator. Trying to interpret the premise that way leads to nonsense, since there would be no grammatical sentence left on which to operate ('O' stands for the operator, with the sentence it binds in brackets):

O[names the relation between Lois and the action-type *donate to charity*]

This is not a formula of any language. It's a misuse of the sentential operator, so clearly it won't do. The premise is therefore purely descriptive. And the conclusion is purely normative, since there is an active normative operator binding the whole sentence:

O[Lois donates to charity]

So although Schurz's view differs from NHG, it fails for the same reasons.

The failure of the narrow approach drives us to the wide one. It has no trouble handling the D-argument, because it designates as a normative sentence any premise that uses or mentions a normative term. Applying it to Pigden's PHG gives us this:

WHG: No argument from premises that do not use or mention normative terms to a conclusion that uses or mentions normative terms is valid, unless the normative terms in the conclusion are vacuous.

This principle is much more secure. The D-argument can't be a counter-example because it mentions 'ought', and Pigden has a proof of it. The issue with WHG is not falsity; it is irrelevance. If WHG is going to be the Guillotine, the wide term approach to normative sentences must be extensionally adequate. But consider the following:

(5L) The word 'ought' has five letters in it.

The wide approach incorrectly calls this a normative sentence. WHG is true, but pairing it with the wide approach makes it a Guillotine for a language that no one speaks. In English, 5L is just a piece of orthography—even if it's normative in some sense, it's not the sort of moral claim about which defenders of the principle care. If merely describing a norm like the spelling

⁹See Schurz [2010a: 204] for discussion.

of 'ought' makes a sentence normative for the purposes of the Guillotine, then Prior's dilemma wouldn't have been a problem in the first place, since the tea-drinking claim describes a norm too.

Some might counter that admitting too much into the set of normative sentences isn't a problem. As long as WHG governs all the sentences we need it to govern, maybe it's acceptable for it to govern some that weren't intended. I agree that it's not a problem if there's reason to group the unintended cases with the original ones. To draw an analogy, it doesn't count against laws meant to regulate telegraph transmissions that they also rule over telephone and internet traffic. But in this case there are good reasons against grouping sentences like 5L with the normative ones. Grammar and linguistics are not ethics, nor anything like it. To see this, notice that being a sceptic about ethics gives us no reason to doubt grammar.

Thus we should conclude that the term approach will not deliver Hume's Guillotine.

4: Second Approach: Relativity

A different approach to Prior's dilemma is to relativize the Guillotine such that neither of his arguments can be sound counter-examples. Toomas Karmo [1988] suggests that a sentence is normative or descriptive relative to a world. Worlds are composed of three maximal consistent sets: a set of obviously normative sentences (like 'Lois ought to donate to charity'), a set of obviously descriptive ones (like 'snow is white'), and a set of non-obvious sentences (like 'snow is white or Lois ought to donate to charity'). A non-obvious sentence is normative just in case changing the truth-value of some obviously normative one, while holding the obviously descriptive ones fixed, changes its truth-value. So, for a non-obvious case like 'snow is white or Lois ought to donate to charity', its status will depend on the world of evaluation. If snow is actually white, then varying obviously normative sentences won't change the truth-value because the first disjunct is fixed. But if snow isn't white at the world of evaluation, then the truth-value depends on the obviously normative sentences because the second disjunct is in that set. Call this the *relativity approach* to normative sentences.¹⁰

This view allows Karmo to prove a candidate for the Guillotine that avoids Prior's dilemma [Karmo 1988: 256].

KHG: There is no argument whose premises are descriptive and whose conclusion is normative, relative to worlds where it's sound.

Given the relativity approach, KHG avoids Prior's dilemma because the crucial tea-drinking disjunction varies between the two arguments. It's descriptive as the conclusion of A because its truth is preserved by the premise. It's true, regardless of what ought to happen to New Zealanders. As a premise in B it's normative. Its truth depends on normative facts about Kiwis because B is sound only if tea-drinking is scarce.

¹⁰Karmo's view draws on Humberstone [1982], which is inspired in turn by Shorter's [1961] reply to Prior.

Some normative-sounding assertions create a problem for the relativity approach. Consider the following:

(BILL) Bill was right to tell the truth about Monica.

To know whether BILL is normative or instead descriptive, we need to specify a world. Start with a world where Bill lies.

Since BILL is a factive claim about satisfying a norm, it's equivalent to a conjunction: *Bill told the truth about Monica and Bill ought to tell the truth about Monica*. We'll run Karmo's test on this for simplicity. If varying an obviously descriptive sentence while fixing the descriptive ones forces it to change truth-value, we know it's normative. But since its first conjunct is descriptive and false, that won't happen. The conjunction is descriptive when Bill lies, so BILL is too.

That BILL is descriptive at any world where he lies already invites a charge of irrelevance, because it means that KHG is a rule for the wrong object language. It leaves out sentences that should be in the scope of the view. And the divergence from natural language gets worse. Notice that while BILL is descriptive at the actual world, it is normative according to the view at any world where Bill does tell the truth. Then consider:

(BILL*) Bill was right *not* to tell the truth about Monica.

Parallel reasons show that BILL* is normative anywhere that Bill doesn't tell the truth. But there is no normatively significant difference between the two—each makes a clear moral evaluation. The only difference is that at some worlds the sentences correctly describe Bill's action and in others they don't.¹¹ Yet it's difficult to see how this would be relevant to assessing normativity. If it isn't relevant, Karmo's approach doesn't accurately model natural language.

Some might argue that correctly describing Bill's action is normatively relevant, by comparing these cases with Prior's disjunction. Because the disjunction would be descriptive when it describes the facts about tea-drinking correctly, and normative when it doesn't, it gets a mixed treatment, like BILL and BILL*. If it's acceptable for Prior's disjunction to vary with correctness, then perhaps it really is relevant to whether a sentence is normative.

Yet even if we accept the mixed treatment of Prior's disjunction—and we needn't—that would show that correct description is normatively relevant only if correctness does some work toward explaining why we accept different verdicts. Otherwise, correctness might have nothing to do with normativity. Other explanations are plausible: the mixed treatment of Prior's disjunction is tolerable because of what asserting it would commit us to at

¹¹I've heard it suggested that sentences with presupposition failure lack truth-values. I am sceptical that this is the case with BILL, but responding in this way won't help HGK. First, if BILL has no truth-value, it cannot be normative on Karmo's definition. Second, since BILL appears to imply truth-apt sentences like 'Bill told the truth', there's a further problem of explaining how we can get a sentence with a truth-value out of a sentence without one. People interested in putting up inference barriers might not want to allow this. Thanks to an anonymous referee for the second point.

different worlds. At worlds where we know that tea-drinking is common in England, we can assert the disjunction while denying that New Zealanders ought to be shot. But when we consider worlds where we know that tea-drinking is not common, asserting the disjunction commits us to saying that all New Zealanders should in fact be shot. Karmo's relativity approach reflects the fact that at some worlds we would be committed to obviously normative claims, but not so at other worlds.

Notice there is no similar change in our commitments when we assert BILL or BILL*. Whatever the world, saying that Bill was right to tell the truth about Monica means that Bill ought to tell the truth about Monica. That's a reason for thinking at least some normative sentences stay that way across worlds. If that's right, then KHG governs the wrong language. So it, too, is a poor candidate to be Hume's Guillotine.

5: Third Strategy: Fragility

Perhaps a different modal logic is the answer. Deontic logic is one where the concepts of permission and obligation are defined like possibility and necessity operators, but within a special set of morally satisfactory worlds. Roughly, a sentence is permissible just when it's true at least one morally satisfactory world, and a sentence is obligatory just when it's true at all the morally satisfactory worlds.

In a deontic system, the truth of ought-sentences is sensitive to changes to the set of worlds that are morally satisfactory in a given model. Greg Restall and Gillian Russell use this feature to dodge Prior's dilemma. Call it the *fragility approach*. It designates a sentence as normative just in case there is at least one model where replacements and additions to the set of satisfactory worlds change its truth-value—that is, one model in which it's fragile—and descriptive just in case it isn't sensitive to those replacements in every model where it's true, making it preserved [Restall and Russell 2010: 254–5]. A sentence like 'Lois ought to donate to charity' is normative on this approach because its truth does change with the set of satisfactory worlds: if Lois donates at all the satisfactory worlds, and we add to the model a new satisfactory world in which she doesn't, the sentence goes from true to false. Conversely, a sentence like 'snow is white' comes out descriptive, because it doesn't matter what's in the set of satisfactory worlds. Since 'snow is white' is not a deontic modal claim, its truth-value doesn't vary with things happening at other possible worlds.

This approach is powerful enough that Restall and Russell don't need to state an alternative to Naïve HG. It's left untouched by Prior's arguments because the crucial tea-drinking disjunction is neither normative nor descriptive. It fails the test for normativity in any model where tea-drinking is actually common in England, because it will remain true no matter what additions or replacements are made to the set of satisfactory worlds. So it's not fragile. If tea-drinking is actually scarce in England, then it fails the test for being descriptive. The disjunction would be true only if all New Zealanders are shot at every satisfactory world. Replacing one of these

worlds with one at which they are not shot makes the disjunction false, so the disjunction is not preserved when true. Thus neither of Prior's arguments counts as an is-ought derivation.

Peter Vranas argues that the fragility approach miscategorizes some normative sentences, making the view extensionally inadequate [2010: 264–5]. Consider LC:

(LC) Lois should donate to charity if she is able.

Vranas would say LC is equivalent to *either Lois is unable to donate or she ought to donate to charity*, which is labelled as non-normative by the fragility approach for the same reason as Prior's disjunction: It's not fragile in models where Lois is unable to donate. Vranas believes that's implausible.

This is the right line of argument to press against the fragility version of the Guillotine, but it's insufficient for present purposes. I'm suggesting that the approach cannot generate an inference barrier relevant to the actual practice of ethics. Vranas's argument doesn't show this for two reasons: (a) it fails to establish that fragility approach really is inadequate, and (b) as we saw with the wide term approach, inadequacy on its own may not lead to irrelevance. More needs to be said.

The argument doesn't show what it is intended to show, because it says nothing about the scope of the obligation operator in LC. This is crucial. If it has wide scope, as in $\Box(AI \rightarrow Dc)$, then the sentence will not be equivalent to the non-normative disjunction $\neg AI \vee \Box Dc$. A wide-scoped LC would come out as normative because, in any model where it's true, adding or replacing satisfactory worlds can make the sentence false, making it fragile. Just add a world where some people refuse who are able to donate. The fragility approach is inadequate only if we can show that LC has narrow scope: $AI \rightarrow \Box Dc$. Given LC's surface grammar, it's not obvious which reading is correct.¹² So, rather than beginning with a sentence of English and insisting on narrow scope, a different strategy is needed.

A better approach is to start in the meta-language: Take a narrow-scoped formula of deontic logic like $AI \rightarrow \Box Dc$. Then we ask if it has any acceptable natural language translations that are normative. The most straightforward is this:

(LC*) If Lois is able, she ought to donate to charity.

To the extent that the original LC is normative—and it is—LC* is normative, too. That's why the fragility Guillotine is inadequate.

But miscategorizing doesn't always guarantee irrelevance, so another step is needed. Remember the telegraph laws that wound up governing internet traffic. In that case, the wide term approach grouped non-normative sentences in with the normative ones. The fragility approach does the opposite by

¹²Some people, notably John Broome [1999] think that sentences like LC should always be understood to have wide scope. I think that's a non-starter, since narrow-scoped imperatives that look just like LC are easily derivable from $\neg AI$.

putting normative cases in with the non-normative, so it's not a situation where all the sentences that need to be governed are within the scope of the principle. This barrier fails to govern what it should, so it's irrelevant, just as was Jackson's version due to its leaving out valid arguments that aren't formally so.

6: Moore's Law

The challenge created by Prior's dilemma looks difficult to surmount. Each of the term, relativity, and fragility approaches generates a true principle that avoids the counter-examples, but every one of those approaches is alienated from actual normative assertion. This suggests there may be no candidate that is both true and adequate. Philosophers may prove additional candidates that look like the Guillotine; but, given the difficulties above, we should be hesitant to accept them, without thorough scrutiny, as being the genuine article.

The general problem comes from attempting to frame a philosophically significant inference barrier around the distinction between normative and descriptive *sentences*, which is difficult to pin down. Moore's Law steers clear of these problems because it's a semantic barrier: no atomic normative *terms* are synonymous with any atomic descriptive terms, either directly or by substitution. I think Moore's Law can both stand in for the Guillotine and improve on it in an important way. I won't argue for Moore's Law, though I do think the open question argument offers a good reason to accept it. What I will try to show instead is that Moore's Law can deliver a lot of the regulation of ethics that Hume's Guillotine promised.

Of course, Moore's Law is not without its own critics. Some metaethical naturalists think it shows very little. After all, 'the morning star' and 'the evening star' have different meanings and are conceptually independent, but co-refer all the same. So, too, normative terms could be undefinable by descriptive terms, but still refer to the same properties. But this is more a criticism of what we can learn from Moore's Law than it is an indictment of the principle itself. The only view Moore's Law rules out is analytic naturalism: 'goodness' and 'pleasure' can't mean the same thing, for instance. Those who aren't analytic naturalists are free to accept Moore's Law and everything it implies. Which properties are identical, and thus which expressions are co-referring, is another question entirely.

Some may doubt the principle's suitability, for other reasons. If we are called upon to give a precise account of the distinction between normative and descriptive terms, we risk making it a rule for the wrong game, too. Recall, though, that advocates of Hume's Guillotine need to be that careful only because of Prior's dilemma. Had Prior not backed the view into a corner, the barrier could have stood on a more rough and ready version of the distinction. I think the same is true of Moore's Law: provided there is no equivalent problem, it should be enough to specify a list of normative terms—like 'ought', 'should', 'may', 'good', 'right', and 'wrong'—that

contrast with descriptive ones and that leave for another day the question of how to specify the difference between them.

So is there an analogue of Prior's dilemma to make trouble for Moore's Law? Creating one requires choosing a term and asking if it's normative or if it's descriptive. If the Moorean says it's descriptive, then the objector shows that it's synonymous with a normative term. If it's said to be normative, the objector shows it has the same meaning as a descriptive term.

At first glance it's difficult to see how to get this started. In order to follow Prior, we would want to find something disjunctive that has both normative and descriptive elements, like 'quick or wrong'. The trouble is that Moore's barrier is meant to govern only atomic terms, so disjunctive predicates look beyond the scope of the Law. Things aren't so simple, however. Predicate abstraction in logic enables us to represent seemingly compound expressions as atomic predicates of a more precise meta-language.¹³ Phrases like ' x is quick or x is wrong' can be formalized into a unitary constant that takes x as its argument. That is, x satisfies the one-place predicate *quick-or-wrong*, which we understand as Qx —an atomic term.

So far, so good, for the opponent of Moore's Law, allowing that predicate abstraction gives an adequate model for natural language. The next step would be to show that our new single expression means the same as both normative and descriptive terms, just as Prior's tea-drinking disjunction participates in valid implication with normative and descriptive sentences. Here is where the analogy to Prior breaks down. With what would the quick-or-wrong predicate be synonymous? It's doubtful that it could be either 'quick' or 'wrong' on its own, as each represents just a third of the complex. Failing those, no other candidates present themselves. More generally, it's difficult to imagine a clearly normative or clearly descriptive term of natural language with a disjunctive meaning comparable to 'quick-or-wrong' or anything like it. We could introduce a new arbitrary term and stipulate that it has such a meaning; but that won't help, because it creates a new object language. Moore's Law governs only natural languages. Creating a dilemma in an artificial one is irrelevant.

There may be a sense in which 'quick' and 'wrong' are partially synonymous with the complex predicate. Each does capture part of the meaning. But this won't help either, since objecting that some normative terms are partially synonymous with a descriptive term (or *vice versa*) would be like objecting to Hume's Guillotine by suggesting that some descriptive sentence partially implies a normative conclusion because it could be used in an argument for that conclusion. To get a genuine objection, we need complete valid implication and complete synonymy.

Prior had the rules of disjunctive syllogism and addition at his disposal, which allowed him to piece together valid inferences with the tea-drinking sentence and its disjuncts. There is no equivalent way to break up a unitary predicate, since it has no proper parts with which to work. The move that was needed in order to state the dilemma prevents us from making it work. That's a good reason to think there is no analogue to Prior's argument here.

¹³Thanks to an anonymous referee for helpful comments here.

With those worries out of the way, all that remains is to show why friends of the Guillotine can accept Moore's Law as a substitute. Many present-day advocates should already accept Moore's Law. Karmo's relativity principle assumes Moore's Law, because synonymy between normative and descriptive sentences would make it impossible to vary the set of obviously normative sentences while fixing the obviously descriptive sentences. Restall and Russell's fragility barrier also assumes the Law, since it treats descriptive terms as predicates and normative terms as deontic operators. As a rule, atomic predicative sentences imply deontic modal sentences only in certain models. But if some normative terms had the same meaning as descriptive ones, then some atomic sentences would imply deontic modals in every model, and that can't be so. Pigden has also expressed support for Moore's Law, though he thinks the Guillotine doesn't imply anything about it [1989: 128].

Some who defend the Guillotine have an intrinsic interest in it because they believe it's an instance of the conservativeness of logic: you can't get out what you haven't put in. But many more have been motivated by its ability to guide ethics. I think there are doubts about certain interpretations of conservativeness because of cases like the Clark Kent D-argument, but if Moore's Law can regulate ethics in the other ways the Guillotine was meant to do so, this latter demand can be met.

The Guillotine is employed primarily to explain why certain arguments are invalid, like *everyone does x, therefore x is permissible*. Hume himself appealed to the Guillotine to show that inferences like *God favours x, therefore x ought to be the case* and *x is human nature, therefore x is permissible* are invalid [1739: 469]. But Moore's Law can explain this too, even if Hume himself might not have accepted it. These inferences are not formally valid; so, if they were valid, it would be in virtue of meaning. But if Moore's Law is true, there are no normative-descriptive synonyms, so arguments like these fail because there is no chain of substitution of synonyms for synonyms with which to license them. This explanation improves on the one from the Guillotine. It preserves a sense in which you can't get an 'ought' from an 'is', it does more than simply assert that the inference can't be made, and it avoids the problematic sentence-level distinction.

This might seem too quick, though. Moore's Law denies that normative and descriptive terms mean the same, but failure of synonymy doesn't always result in an inference barrier. The word 'uncle' isn't synonymous with 'male', but the fact that Clark is an uncle validly implies that Clark is male.¹⁴ So, perhaps Moore's principle can't block the arguments that the Guillotine was meant to block, after all.

While it's true that some inferences cross gaps in meaning, there is an important difference between the uncle-to-male argument and the ones Hume had in mind. In the uncle-male case, synonymy does secure the conclusion—not directly, but by substitution. To derive 'Clark is male' from 'Clark is an uncle', we replace 'uncle' with 'father's brother', then 'brother' with 'male sibling', and we have the conclusion by simplification. The human-nature-to-permissible argument (and others like it) cannot work this way if Moore was

¹⁴Thanks to an anonymous referee for this suggestion.

right. If it did, there would have to come a point in the substitution chain where a descriptive term was replaced with a normative one on the basis of meaning. Moore's Law says there is a divide that cannot be bridged—the only terms that can stand in for 'human nature' are descriptive ones, and the only ones that can stand in for 'permissible' are normative ones.

Others have relied on the Guillotine to argue that normative beliefs cannot be deductively justified by descriptive beliefs [Sinnott-Armstrong 2000; Huemer 2005]. Moore's Law cannot show this, of course, since beliefs are propositional and since Moore's Law says nothing about sentence-level things. This is for the best, though, because the same problems that make trouble for the inference barrier will afflict a barrier between normative and descriptive beliefs. Yet Moore's Law can show something similar: no beliefs featuring normative terms can be justified by deduction from a set of premises lacking normative terms by substitution of synonyms for synonyms. Moore's barrier shows that the chain of substitution cannot be completed.

Others have tried to leverage the Guillotine to create an explanatory advantage for moral non-cognitivism, the idea that normative assertions lack truth-values. Non-cognitivists have an easy explanation of why Hume's Guillotine would be true—since normative sentences have no truth-values, they cannot be implied by any sentences whatsoever—while cognitivists are supposed to find it difficult to offer a competing explanation.¹⁵ Present purposes don't require taking a stand on whether non-cognitivists would truly have such an advantage; but, to the extent that there would be one, Moore's Law can be employed for the same work. Non-cognitivists can instead explain the failure of synonymy between sentences using normative and descriptive terms by appealing to the former's lack of truth-aptness, and by claiming as before that cognitivists cannot offer a good competitor.

Replacing the Guillotine with Moore's Law has one additional benefit. Many have thought it inconsistent with the popular Kantian principle that *ought* implies *can*. If Kant was right, then '*S* cannot φ ' gives us 'it's not the case that *S* ought to φ .' This appears to generate an is-ought derivation: if you cannot, then you need not. But there is no conflict between Moore and Kant's rule. Moore's Law requires only that 'ought' not be synonymous with any descriptive terms, and this is perfectly consistent with the Kantian principle, as long as we understand the inference from *ought* to *can* as licensed by something other than synonymy. Since all we need is a relation of valid implication for the inference to go through, synonymy isn't needed. The move from *ought* to *can* goes through if ability necessarily supervenes on obligation, such that any two worlds with identical distributions of obligation would have likewise identical distributions of ability. It may be possible in the end to reconcile the Guillotine and the Kantian principle, but doing so demands a lot of explanation and might involve additional meta-ethical commitments.¹⁶ If we substitute Moore's Law for Hume's Guillotine, we can have *ought* implies *can* without the explanatory debts.

¹⁵See Hare [1952], in particular, for this use of the Guillotine.

¹⁶See Bloomfield [2007] for an excellent attempt at this.

There may yet be a way to rescue Hume's Guillotine from Prior's dilemma, but the current prospects are dim, since the troubles with normative and descriptive sentences cast doubt on the enterprise of proving it. If what I've argued about Moore's Law is correct, then the heavy lifting needed to save the Guillotine from the mire of problems that emerged above is unnecessary. Giving it up involves little if any cost and affords at least one important additional benefit. Devotees would therefore do better to focus instead on defending and invoking Moore's Law.¹⁷

Syracuse University

Received: November 2013

Revised: July 2014

References

- Black, M. 1964. The Gap Between 'Is' and 'Should', *The Philosophical Review* 73/2: 165–81.
- Bloomfield, P. 2007. Two Dogmas of Metaethics, *Philosophical Studies* 132/3: 439–66.
- Broome, J. 1999. Normative Requirements, *Ratio* 12/4: 398–419.
- Castañeda, H. 1973. On the Conceptual Autonomy of Morality, *Noûs* 7/1: 67–77.
- Hare, R.M. 1952. *The Language of Morals*, Oxford: Clarendon Press.
- Huemer, M. 2005. *Ethical Intuitionism*, Basingstoke: Palgrave Macmillan.
- Humberstone, I.L. 1982. First Steps in Philosophical Taxonomy, *Canadian Journal of Philosophy* 12/3: 476–8.
- Hume, D. 1739 (1978). *A Treatise of Human Nature*, 2nd edn, ed. P.H. Nidditch, Oxford: Clarendon Press.
- Geach, P.T. 1977. Again the Logic of 'Ought', *Philosophy* 52/202: 473–6.
- Jackson, F. 1974. Defining the Autonomy of Ethics, *The Philosophical Review* 83/1: 88–96.
- Karmo, T. 1988. Some Valid (But No Sound) Arguments Trivially Span the 'Is'–'Ought' Gap, *Mind* 97/386: 252–7.
- Pigden, C.R. 1989. Logic and the Autonomy of Ethics, *Australasian Journal of Philosophy* 67/2: 127–51.
- Pigden, C.R. 2010a. Introduction, in *Hume on 'Is' and 'Ought'*, ed. C.R. Pigden, Basingstoke: Palgrave Macmillan: 1–38.
- Pigden, C.R. 2010b. On the Triviality of Hume's Law: A Reply to Gerhard Schurz, in *Hume on 'Is' and 'Ought'*, ed. C.R. Pigden, Basingstoke: Palgrave Macmillan: 217–38.
- Prior, A.N. 1960. The Autonomy of Ethics, *Australasian Journal of Philosophy* 38/3: 199–206.
- Restall, G. and G. Russell 2010. Barriers to Implication, in *Hume on 'Is' and 'Ought'*, ed. C.R. Pigden, Basingstoke: Palgrave Macmillan: 243–59.
- Schurz, G. 2010a. Non-Trivial Versions of Hume's Is-Ought Thesis, in *Hume on 'Is' and 'Ought'*, ed. C.R. Pigden, Basingstoke: Palgrave Macmillan: 198–216.
- Schurz, G. 2010b. Reply to Pigden, in *Hume on 'Is' and 'Ought'*, ed. C.R. Pigden, Basingstoke: Palgrave Macmillan: 239–42.
- Searle, J. 1964. How to Derive 'Ought' from 'Is' *The Philosophical Review* 73/1: 43–58.
- Shorter, J.M. 1961. Professor Prior on the Autonomy of Ethics, *Australasian Journal of Philosophy* 39/3: 286–7.
- Sinnott-Armstrong, W. 2000. From 'Is' to 'Ought' in Moral Epistemology, *Argumentation* 14/2: 159–74.
- Vranas, P. 2010. Comments on 'Barriers to Implication', in *Hume on 'Is' and 'Ought'*, ed. C.R. Pigden, Basingstoke: Palgrave Macmillan: 260–7.

¹⁷Thanks to Ben Bradley, John Monteleone, Deke Gould, Greg Ray, and three anonymous referees for helpful comments on earlier drafts of this paper.