

Mind the Is-Ought Gap¹

Forthcoming in the *Journal of Philosophy*

Daniel J. Singer

singerd@phil.upenn.edu

Penultimate draft from February 2015 – Please cite only the published version

Abstract:

The is-ought gap is Hume’s claim that we can’t get an ‘ought’ from just ‘is’s. Prior (1960a) showed that its most straight-forward formulation, a staple of introductory philosophy classes, fails. Many authors attempt to resurrect the claim by restricting its domain syntactically or by reformulating it in terms of models of deontic logic. But, those attempts prove to be complex, incomplete, or incorrect. I provide a simple reformulation of the is-ought gap that closely fits Hume’s description of it. My formulation of the gap avoids the proposed counterexamples from Prior and offers a natural explanation of why they seem compelling. Moreover, I show that my formulation of the gap is guaranteed by standard semantic rules of normative terms, which provides a more general reason to accept it.

According to Hume, we shouldn’t be able to learn about how things ought to be just from how things are. The idea is often summed up as “no ought from an is.” This is the putative *is-ought gap*.

A. N. Prior purported to give several counterexamples to the is-ought gap, and there are many proposed modifications and defenses of the gap in response.² I argue for a new way to understand the is-ought gap. Along with Russell and Restall, I claim that

¹The core idea presented here was developed in conversations with Alex Silk. I am extremely grateful to him for his help on this. I also want to thank Elisabeth Camp, Allan Gibbard, James M. Joyce, Sarah Moss, Peter Railton, Chandra Sripada, W. Robert Thomas, and the University of Michigan Graduate Student Working Group for feedback on previous versions of this work.

² See Prior, A. N. 1960. “The Autonomy of Ethics.” *Australasian Journal of Philosophy* 38 (3): 199–206. For responses, see Jackson, Frank. 1974. “Defining the Autonomy of Ethics.” *The Philosophical Review* 83 (1). Duke University Press on behalf of Philosophical Review: pp. 88–96, Pigden, Charles R. 1989. “Logic and the Autonomy of Ethics.” *Australasian Journal of Philosophy* 67: 127–51, and Russell, Gillian, and Restall, Greg. 2010. “Barriers to Implication.” In *Hume on Is and Ought*, edited by Charles Pigden. Palgrave MacMillan.

the is-ought gap is best understood in terms of the semantics, rather than the syntax, of arguments.³ Unlike Russell and Restall though, I provide a complete and correct formulation of the is-ought gap that offers a natural explanation of why the counterexamples fail and stays true to Hume's original motivation.

Introducing the Gap

In Book III, Part I, Section I of the *Treatise of Human Nature*, Hume gives the first statement of the is-ought gap. He says:

In every system of morality, which I have hitherto met with, I have always remark'd, that the author proceeds for some time in the ordinary ways of reasoning, and establishes the being of a God, or makes observations concerning human affairs; when of a sudden I am surpriz'd to find, that instead of the usual copulations of propositions, *is*, and *is not*, I meet with no proposition that is not connected with an *ought*, or an *ought not*. This change is imperceptible; but is however, of the last consequence. For as this *ought*, or *ought not*, expresses some new relation or affirmation, 'tis necessary that it shou'd be observ'd and explain'd; and at the same time that a reason should be given, for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it.⁴

Hume has us imagine a hypothetical interlocutor seeking to give a theory of morality. The interlocutor offers an argument from wholly non-normative premises to a conclusion about how things ought to be. This inference of a normative claim from only descriptive facts, Hume seems to be claiming, is problematic because "this new relation" of normativity seems "entirely different" from descriptive or observational propositions. What *is* the case, the observational or descriptive facts, do not seem to determine at all how things *ought* to be. This is a rough first pass at the is-ought gap:

SIMPLE GAP No normative truth is determined by any non-normative truths.⁵

SIMPLE GAP, if right, would have it that no collection of the descriptive facts is enough to establish how things ought to be. SIMPLE GAP, as stated, is ambiguous between an *epistemic* reading and a *metaphysical* reading. On the epistemic reading, the claim is

³ Op. cit.

⁴ Hume, David. 1978. *A Treatise of Human Nature*. Edited by L. A. Selby-Bigge and P. H. Nidditch. 2nd ed. Clarendon Press ; Oxford University Press, Oxford : New York. P. 469

⁵I formulate the simple version of the claim here in terms of normative and non-normative *truths*. It is thus formulated to mirror Hume's talk of propositions, which I take to be bearers of truth-values, though the use of 'truth' in the claim is unnecessary. If for example, one denies the possibility of normative truths, one may substitute whatever analogue of truth one wishes here or simply formulate the claim without any such analogue: "Nothing normative is determined by anything non-normative" (though I take this formulation to be less clear).

about what we can properly *infer*: We cannot draw any conclusions about what ought to be the case just from what is the case. Such an argument, according to the epistemic reading, would be invalid.

Since Hume uses the gap to reject the arguments of his contemporaries, Hume seems to accept the epistemic reading of the claim. But, Hume doesn't defend it by pointing to any formal, syntactic, or conceptual problems of their arguments. Hume's defense of the gap is metaphysical: it's about whether there really is that connection between the descriptive and the normative that the offending arguments need. The normative conclusions "[express] some new relation or affirmation, 'tis necessary that it shou'd be observ'd and explain'd; and at the same time that a reason should be given; for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it," he tells us.⁶

The metaphysical reading of SIMPLE GAP expresses this deeper disconnection between the normative and the non-normative. On the metaphysical reading, what is the case normatively isn't fixed by non-normative facts. So, if there were a creator who set up all of the truths, the metaphysical version of the is-ought gap would require that the creator to set the normative truths even after all of the non-normative truths were established.

As Hume's argument requires, there is a natural connection between the two versions of the gap claim. If the metaphysical claim is true, the epistemic claim would be true as well: one could not properly argue from just non-normative facts to any normative facts (without any bridge principles as premises). This is because there would be some possibilities not ruled out by the premises that are ruled out by the conclusion, so such an argument would be invalid. Following Hume, I'll restrict my focus here to the epistemic reading of the claim, the claim that *we cannot properly reason* from claims about how things are to a claim about how things ought to be.⁷

⁶ Op. cit. p. 468

⁷For more on the metaphysical version of the claim, see Blackburn, S. 1971. "Moral Realism." In *Morality and Moral Reasoning*, edited by John Casey, 101–24. London, Methuen for a discussion of the supervenience of the normative on the non-normative and 2013. "Autonomy of Ethics." In *International Encyclopedia of Ethics*. Blackwell Publishing Ltd for remarks on how this affects the metaphysical claim.

Problems with the Simple Conception of the Gap

A. N. Prior showed that the epistemic reading of SIMPLE GAP is false with three counterexamples:⁸

(A)

1. Tea-drinking is common in England.

2. Therefore, either tea-drinking is common in England or all New Zealanders ought to be shot.

(B)

1. There is no man over 20ft high.

2. Therefore, there is no man over 20ft high who ought to sit in an ordinary chair.

(C)

1. Undertakers are church officers.

2. Therefore, if church officers ought to be reverent, undertakers ought to be reverent.

Each example consists of a formally valid argument with non-normative premises and a normative conclusion. They show that the epistemic reading of SIMPLE GAP is false: we *can* properly reason from non-normative premises to a normative conclusion.⁹

⁸ See Prior, *op. cit.* I provide slightly modified versions of the examples from Pigden, *op. cit.* p. 132, of each of Prior's examples.

⁹Though Prior did not think of a claim like SIMPLE GAP as ambiguous between the two given readings, the counterexamples also show against the metaphysical reading of the claim: since there are valid arguments from non-normative facts to some normative facts, that means that some non-normative facts must fix some normative facts.

Prior also noted that there seems to be something deficient about each of the conclusions of these argument, in that they don't seem to be action-guiding or *genuinely* normative. On that line, one might attempt to avoid the worries by identifying that deficiency and omitting similarly deficient conclusions from the intended scope of the is-ought gap.

Prior considers this move, but he argues that it won't work by showing how each of the conclusions can be used in the context where it does have interesting, possibly

Some Attempts to Avoid the Counterexamples

Several authors have proposed responses to Prior's worries for the is-ought gap. A popular move here is to limit the intended domain of the is-ought gap to arguments with sufficiently simple normative conclusions. The underlying idea is that none of Prior's examples are counterexamples to the is-ought gap because none of them entail a simple statement of what ought to be the case or what we have reason to do. Gibbard, for example, argues that it's commonplace for hypothetical ought claims to follow from descriptive premises, and he offers a restriction of the claim to sufficiently simple conclusions, where a normative conclusion is *simple* when its main connective is Gibbard's primitive subjective ought.¹⁰

But, restricting the gap to only arguments with simple conclusions is both too strong and too weak. It's too strong because it excludes some arguments from the purview of the is-ought gap that it should not. For instance, it excludes an argument with the conclusion "Either Jane ought to eat tomato soup, or Ange ought to buy garlic bread." This conclusion isn't a simple normative sentence in Gibbard's sense. But that conclusion tells us something substantive about what ought to be the case, and the is-ought gap, properly formulated, ought to deny the validity of arguments with only non-normative premises and that conclusion. (Note that to show that Gibbard's restriction is too strong, I need not produce a valid argument from descriptive premises to this conclusion. Doing so would be to pronounce the fate of the is-ought gap, not merely clarify its intended domain. Rather, to show Gibbard's restriction to be too strong, I must only show that *if there were such an argument*, it would be a counterexample to the is-ought gap. An argument with only non-normative premises and the given disjunctive conclusion would be one that would "[express]

action-guiding, normative force. For example, we can get the conclusion of the first argument like this:

1. Anyone who does what is not common in England ought to be shot;
 2. All New Zealanders drink tea;
-
3. Therefore either tea-drinking is common in England or all New Zealanders ought to be shot.

This example shows that, contra the intuition that there's no real force behind the normative terms in the conclusion of the first argument, it is not always normatively impotent. So, we can't merely restrict the class of conclusions that the is-ought gap is meant to apply to. Prior concludes then that there is no gap.

¹⁰See Gibbard, Allan. 2012. *Meaning and Normativity*. Oxford University Press. p. 80-81. Note that Gibbard does take himself to be responding to Prior here. Rather, he is considering some potential outs for a puzzle about correct belief. He does conclude though that "we can maintain ...that no elementary *ought* follows analytically from an *is*" (op. cit. p. 88), in light of his consideration of that puzzle.

some new relation or affirmation” for which “’tis necessary that it shou’d be observ’d and explain’d.”)

Thinking of the is-ought gap as only applying to arguments with simple normative conclusions is also too weak. Gillian Russell attributes the following example to Gideon Rosen:¹¹ Suppose “to flurg” means to do something that one ought not do in front of children. Then we could deduce from “Lauren is in front of children.” to “Lauren ought not flurg.” The conclusion here is simple, in Gibbard’s sense, but it also follows from the non-normative premise.¹² So, modifying the domain of the gap claim to make it apply only to simple normative sentences is too strong in that it doesn’t capture all of the intended domain of the is-ought gap, and it is too weak in that it doesn’t save the claim of Prior-style counterexamples like this one from Russell.

Other authors have proposed modifications of the is-ought gap claim that, like Gibbard, attempt save it by limiting its intended domain on the basis of the syntax of the arguments. In Prior’s original essay about the problematic arguments, for example, he notes that the normative terms in his first two examples appear not to do any normative work; in them, “the duty established is not one that we need ever be practically anxious about.”¹³ Prior diagnoses this as what he calls “contingent vacuousness”. A term is contingently vacuous when it can be replaced with a grammatical counterpart without sacrificing the validity of the argument. Prior’s first two examples would be still be valid if ‘ought’ were replaced by ‘want’, for example. Prior fears that this cannot be the feature that sets his counterexamples to the is-ought gap apart from other arguments though, because he thinks that the ‘ought’ in his third example is not contingently vacuous.¹⁴

Frank Jackson also considers examples like Prior’s and diagnoses them as being being *ethically invariant*, a property that is formally similar to Prior’s contingent

¹¹ Gillian Russell. 2010. “In Defence of Hume’s Law.” In *Hume on Is and Ought*, edited by Charles Pigden. Palgrave MacMillan. P. 145.

¹²One might worry that stipulating concepts like FLURG is illegitimate and subject to worries like those surrounding the concept TONK as discussed by Prior, Arthur N. 1960. “The Runabout Inference Ticket.” *Analysis* 21: 38–39. In general, that seems right, but this particular case is not worrisome. FLURG is plausibly quite similar to everyday thick normative concepts like OBNOXIOUS: being obnoxious requires satisfying some descriptive criteria, such as being loud or unrelenting, as well as being such that one ought not be that way. The example of ‘flurg’ is used since it’s easier to specify in the stipulated case what exactly the descriptive content is.

¹³ Prior, A. N. 1960. “The Autonomy of Ethics.” *Australasian Journal of Philosophy* 38 (3): p. 203

¹⁴ op. cit. p. 204-6.

vacuousness.¹⁵ To avoid other potential counterexamples to the is-ought gap, Jackson proposes an intricate formulation of its domain, claiming that it permits arguments with normative conclusions only if they are factual-invariant, ethical-invariant, or reducible to such by synonymy substitution.¹⁶

Pidgen, following MacIntyre, objects to Jackson's method by suggesting that without motivating such a conception of the is-ought gap by means of a more general principle, we can expect nothing more than the mere, "perhaps temporary, drying up of counterexamples."¹⁷ Pidgen then attempts to provide a more general motivation for the is-ought gap, which treats it as an instance of the claim that logic is conservative, i.e. that nothing substantive can occur in the conclusion of a deductively valid argument that is not somehow contained in the premises. Pidgen says that we should treat a logic as conservative when "A predicate or propositional variable cannot occur non-vacuously in the conclusion of a valid inference unless it appears among the premises." Then, the conservativeness of our standard logics vindicates the is-ought gap, according to Pidgen, since it shows us that we cannot get an 'ought' claim in the conclusion of a valid argument with 'ought'-free premises unless that 'ought' is vacuous.¹⁸

Pidgen's proposal, like the other proposed solutions I discuss above, construes the is-ought gap as essentially making a claim about when and how certain words like 'ought' can appear in an argument. As such, those proposals treat the gap as primarily a syntactic or logical claim. This spells trouble for Pidgen's approach, as he notices, since he must deny the possibility of an illuminating deontic logic, since such a logic would treat some normative claims as non-vacuously introducible into a conclusion of a valid argument.¹⁹ I hope we can accept the is-ought gap without

¹⁵ See Jackson, Frank. 1974. "Defining the Autonomy of Ethics." *The Philosophical Review* 83 (1). Duke University Press on behalf of Philosophical Review: pp. 88–96. Formally, an argument is ethically invariant when the normative terms in it can be replaced uniformly with any other grammatically appropriate term without sacrificing argument validity (p. 91). This is a slightly different property than Prior's contingent vacuousness; see Pidgen, Charles R. 1989. "Logic and the Autonomy of Ethics." *Australasian Journal of Philosophy* 67: 133-4 for an explanation of the difference and discussion of it. When an argument is ethically invariant, its validity is, in a sense, independent of the ethical terms.

¹⁶An argument is Σ -invariant when its validity is independent of the meaning of the terms of type Σ . See op. cit. p. 91-3.

¹⁷ See Pidgen, op. cit. p. 134-5 and MacIntyre, Alasdair C. 1981. *After Virtue : A Study in Moral Theory*. University of Notre Dame Press. p. 54-5.

¹⁸ Pidgen, Charles R. 1989. "Logic and the Autonomy of Ethics." *Australasian Journal of Philosophy* 67: 127–51.

¹⁹This is because Pidgen cannot treat 'ought' as a sentential operator, like deontic logics do, because there would be truths of deontic logic that contain 'ought' non-

denying the possibility of an informative deontic logic, especially given the progress that has been made by treating normative terms as modal terms in linguistics.²⁰ So following Pigden's and MacIntyre's charge, in the next section, I provide a formulation of the is-ought gap that flows from a more general understanding of the idea. I avoid the problems of the approaches I mention above by treating the is-ought gap, not as a syntactic or logical claim, but as a semantic claim.

A New Formulation of the Is-Ought Gap

For a more general defense of the gap, we should look more closely at Hume's problematic case of the moral theorist who gives an argument with solely non-normative premises to a normative conclusion. Notice that when we imagine this case, we don't imagine that the conclusion is just *any* normative claim; since she is a moral theorist, we imagine Hume's interlocutor to be defending a claim about what ought (or ought not) to be the case in the actual world. For this reason, the moral theorist would not be offering an argument like Prior's proposed counterexamples. That's because the conclusions of those arguments don't offer any substantive normative claims *to someone who accepts the premises*: Consider argument (A), the argument from 'Tea-drinking is common in England' to 'Either tea-drinking is common in England or all New Zealanders ought to be shot.' This would be an odd argument for a real moral theorist to offer. But, its oddness can't be due to the argument being invalid (because it isn't). Rather, what's odd about the argument is that the normative impact of the conclusion is irrelevant to the possibilities being reasoned about. If the premises are true, the conclusion doesn't tell us anything about how things ought to be *in the actual world*. For that reason, (A) would not be the sort of argument Hume's moral theorist would give in defending his moral theory. More generally, the point is that when our goal is to find out what ought to be the case, like we do in moral theorizing, arguments to normative conclusions are useful only insofar as those arguments offer substantive and relevant normative constraints on the possibilities under consideration.

vacuously and that can be deduced from no premises (and so any 'ought'-free sentences as well), counter to the way he employs the conservativeness of logic. Perhaps $O\phi \rightarrow \neg O\neg\phi$ (If it ought to be the case that ϕ , then it is permissible for it to be the case that ϕ) is one such example that can be deduced from any premises.

For another critique of Pigden, see Nelson, Mark T. 1995. "Is It Always Fallacious to Derive Values from Facts?" *Argumentation* 9 (4). Springer Netherlands: 553–62.

²⁰See, for example, the now-standard analysis of normative terms as modal terms from Angelika Kratzer. See Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1: 337–55 as well as Kratzer, Angelika. 1991. "Modality." In *Semantics: An International Handbook of Contemporary Research*, edited by A. von Stechow and D. Wunderlich, 639–50. Berlin, Germany: de Gruyter.

When we introduce and motivate the is-ought gap claim in teaching contexts, we also only consider arguments where the normativity in the conclusion is substantive and relevant to the possibilities described by the premises. In an introductory philosophy class, for example, we might try to convince students that there is such a gap by giving an example of fallacious argument offered by disputants in a fight over a real moral issue. For example, we might introduce the argument from “Abortion kills babies” to “Abortion is wrong” and expect the students to have the intuition that the argument is invalid. We would not try to motivate the gap claim with an argument where the conclusion was irrelevant to the premise, such as the argument from “Abortion kills babies” to “If abortion doesn’t kill babies, it is wrong.” The best explanation of this is that the gap claim is really the idea that *we can’t advance our normative inquiry* by considering an argument from non-normative premises to a normative conclusion, not that there are no valid arguments of that form.

My proposed formulation of the is-ought gap takes these considerations seriously by restricting the domain of is-ought gap to arguments where the normativity in the conclusion is relevant to the possibilities described by the premises.

IS-UGHT GAP There are no valid arguments from non-normative premises to a relevantly normative conclusion.

where a conclusion of an argument is *relevantly normative* when it has substantive normative implications for the possibilities described by the premises (assuming there are some such possibilities). A more fleshed out account of when a conclusion counts as relevantly normative is still required, and below I’ll give a natural one.

Before that though, notice that none of Prior’s proposed counterexamples show against the IS-UGHT GAP formulated thusly. Each of those arguments has the property that if the premises are true in the some world, then the conclusion offers no substantive claim about what ought to be done in that world. Consider argument (A) again. If we have two moral theorists who both agree that tea-drinking is common in England but disagree about what ought to be the case, such as about whether New Zealanders ought to be shot. Learning the conclusion of (A) surely won’t help them reach consensus. The conclusion is only substantively normative in possibilities they aren’t considering, namely ones in which tea-drinking isn’t common in England. A similar lesson applies to Prior’s other examples, so none of those threaten IS-UGHT GAP.

The Is-Ought Gap in Normative Semantics

Hume gave us an intuitive motivation for IS-UGHT GAP. Here I take the case one step further by showing that IS-UGHT GAP, when properly formalized, should be seen as a theorem of normative semantics. If that’s correct, the is-ought gap is not subject to Prior’s or any other counterexamples. To show this, I assume that normative sentences/utterances are interpreted with respect to points of evaluation that consist of (perhaps among other things) an ordinary possible world and a normative standard. This is an extremely minimal assumption that allows me to remain

agnostic about interpreting the normative semantics as relativist, contextualist, or invariantist.

To see just how minimal the assumption is, recall that on the double indexing sort of semantics inherited from Stalnaker, Kaplan, and Lewis, points of evaluation are structured in terms of a context that takes expressions into contents, and an index (or circumstance of evaluation) that takes contents into extensions. This kind of semantics is very general and is meant to provide a foundation for doing semantics of all kinds of language. The double indexing semantics is often supplemented with different elements, like times, places, worlds, and individuals to account for different aspects and types of language, like temporality and modality. To account for normative language, the double indexing picture is supplemented with a normative standard. According to the contextualist about normative language, the normative standard figures in the context and affects what proposition is expressed. According to the relativist, the normative standard figures into the index and affects the truth value of the proposition expressed in a context. According to the invariantist, the normative standard is contained in the structure of the possible world, which is itself standardly treated as a coordinate of the index. So, each position admits the very minimal apparatus that I need, that sentences are evaluated with respect to a possible world and a normative standard.²¹

Here I will show that IS-UGHT GAP, properly conceived of, is a theorem of a particular normative semantics, the one given by Gibbard's *Thinking How to Live*.²² This particular semantics makes it easy to formalize IS-UGHT GAP, show how it is a theorem, and see how to generalize these lessons to other semantics of normative language that meet the minimal assumption. Gibbard uses his semantics in a defense of an expressivist account of normative language, but my use of the framework is completely orthogonal to that issue.

On Gibbard's semantics, the contents of sentences and utterances are modeled as sets of pairs, each consisting of a possible world and a complete plan for action in any hypothetically possible scenario. Then, like in a standard possible-world semantics, we model the content of a sentence 'P' with the set of world-norm pairs that are compatible with P. The contents of non-normative sentences like 'Snow is white' are represented as the collection of world-norm pairs where the world is one in which snow is white. The semantics of 'snow is white, or Mark ought to cook' is

²¹ The three positions simply disagree about how best to account for those two aspects of the semantics. Also, in order to distinguish the interpretations of normative sentences that differ in truth value, the invariantist will either need to allow impossible worlds that differ in which normative standard is contained in the world or allow that the normative aspects of worlds are not fixed by the non-normative aspects. Many thanks to Alex Silk and Elizabeth Camp for helping me to clarify the flexibility of my proposal in this respect.

²² Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge: Harvard University Press.

given by the set of world-plan pairs where either snow is white in the world or the plan calls for Mark cooking.

A sentence will be non-normative, on this kind of semantics, when it is norm-invariant, i.e. if each world in the set is paired with every possible plan. More explicitly, if W is the collection of all possible worlds and N is the collection of all complete plans for action, a set of world-norm pairs $P = \{\langle w_i, n_j \rangle\}$ is *norm-invariant* iff $(\forall w \in W)(\exists n(\langle w, n \rangle \in P) \rightarrow (\forall n \in N)(\langle w, n \rangle \in P))$.²³ A set of world-norm pairs that is norm-invariant is compatible with all possible plans. Such a set therefore doesn't exclude any normative possibilities, so the sentence or utterance that it represents must be non-normative.

Before considering IS-UGHT GAP in this semantics, consider first how to formalize SIMPLE GAP: Take an argument from premises $\{P_i\}$ to conclusion C . SIMPLE GAP would say this: If each of the P_i is norm-invariant, then C is norm-invariant. As we should expect, Prior's (A) shows against this formulation of the claim. The premise, that tea-drinking is common in England, is norm-invariant, but the conclusion, that either tea-drinking is common in England or all New Zealanders ought to be shot, is not. The set of world-norm pairs that represents the conclusion includes a pair where there is no tea and all New Zealanders ought to be shot, but it doesn't include a tea-less world where some New Zealanders ought not to be shot. So, it is not norm-invariant, and this world-norm-style analogue of SIMPLE GAP fails, as expected.

The problem I diagnosed with Prior's counterexamples is that they only make substantive normative claims about worlds that are incompatible with the premises: the normative aspects of the conclusion doesn't tell us how things ought to be in the worlds the argument is about. The solution then is to restrict the domain of the is-ought gap to arguments in which the normative aspect of the conclusion is relevant to the possibilities being reasoned about. We can formalize this intuition in our semantic framework easily. To decide whether the conclusion of an argument makes a claim about how things ought to be in the worlds described by the premises, we decide whether the conclusion is norm-invariant when restricted only to the worlds compatible with the premises. In our semantics, when the premises are norm-invariant, deciding this is equivalent to deciding whether the conjunction of the conclusion and the premises is norm-invariant. This then is a reformulated version of IS-UGHT GAP in Gibbard's semantics:

WORLD-NORM GAP If $\{P_i\} \vdash C$, each of $\{P_i\}$ is norm-invariant, and $P_1 \wedge P_2 \wedge \dots$ is satisfiable, then $P_1 \wedge P_2 \wedge \dots \wedge C$ is norm-invariant.

Intuitively, WORLD-NORM GAP tells us that if the premises of an argument are norm-invariant, then the set of all world-norm pairs compatible with the conclusion *and the premises* is also norm-invariant. By checking the conclusion conjoined with the premises for norm-invariance, we restrict our attention to only those worlds where

²³This is a restatement of the definition of 'norm-invariant' employed by Schroeder, Mark. 2011. "Attitudes and Epistemics." (Manuscript).

the premises are true. The condition that the premises be satisfiable is included to avoid the special case where non-norm-invariant claims follow trivially from contradictory premises.

As I claimed, WORLD-NORM GAP avoids Prior's counterexamples. Consider this one:

(C)

1. Undertakers are church officers.
2. Therefore, if church officers ought to be reverent, undertakers ought to be reverent.

Since the argument is valid, the conclusion must be satisfied by at least the same world-norm pairs that satisfy the premise. Therefore, once we conjoin the conclusion with the premise, we're left with a claim that is logically equivalent to the premise. Since the premise is norm-invariant, it follows that the conclusion conjoined with the premises is also norm-invariant. Hence, that argument is compatible with this formulation of the is-ought gap. Prior's other supposed counterexamples are handled similarly.²⁴

One might worry about avoiding the counterexamples to SIMPLE GAP by limiting the scope of the is-ought gap to arguments from premises to the conjunction of those premises and the conclusion. In doing so, it ignores the possibility that normativity may arise in a conclusion that is strictly weaker than the conjunction of the premises.

But this is feature, not a bug. If normativity arises only when the conclusion of an argument is strictly weaker than the conjunction of the premises, no reasoner could use that argument to come to a conclusion about what they ought to do. The world-norm semantics gives us a natural framework for understanding why this is so: Suppose Rachel wants to know whether she ought to donate to charity. We can represent Rachel's mental state by the collection of world-norm pairs M compatible with what she believes and the norms she accepts. For Rachel to become decided on whether she ought to donate, it is for the plan in each world-norm pair in M to call for donating in the (centered) world of the pair. So, if Rachel isn't already committed on whether to donate, there must be two world-norm pairs $\langle w_1, n_1 \rangle$ and $\langle w_2, n_2 \rangle$ in M such that n_1 calls for donating in w_1 and n_2 permits not donating in w_2 . Now suppose Rachel happens across a valid argument where the premises are norm-invariant but the conclusion is not. The only way for Rachel to become committed to donating in response to that argument is for her to change her mental state so that one of those two world-norm pairs is removed from its representation. Could Rachel do this as a proper response to a valid argument? No: Since the argument is valid,

²⁴E.g. the conjunction of 'There is no man over 20ft high' with the conclusion that 'There is no man over 20ft high who ought to sit in an ordinary chair' is first-order equivalent to 'There is no man over 20ft high,' which is norm-invariant.

Rachel must either accept the conclusion or reject a premise. The premises are norm-invariant, so we can think of them as ordinary centered possible world propositions. Rachel accepts them when every world in a world-norm pair in M is contained in every premise. So, if Rachel does accept them, w_1 and w_2 must be in every premise. Since the premises are norm-invariant, the premises must also have w_1 and w_2 paired with every possible norm, so importantly, $\langle w_1, n_1 \rangle$ and $\langle w_2, n_2 \rangle$ must be in every premise. But, since those world-norm pairs satisfy the premises and the argument is valid, they must also satisfy the conclusion. So when Rachel accepts the conclusion, those pairs will remain in her mental state, and she'll remain uncommitted about whether to donate. On the other hand, if Rachel does not accept the premises of the argument, her mental state will not change either. So again, she will remain uncommitted about whether to donate.²⁵

The key claim of IS-UGHT GAP is this: for arguments from non-normative premises to a normative conclusion, none of the genuinely normative aspects of the conclusion can be relevant to the possibilities described by the premises. But, since a deductive argument could only help us learn something about how things ought to be insofar as we accept the premises, any potential normative guidance that could be derived from non-normative premises must only apply in possibilities where the premises fail. When we ignore those possibilities by restricting our attention only to

²⁵This proof works for deductive arguments because deduction is non-ampliative. By accepting the conclusion of such an argument, the reasoner is not committed to anything she wasn't committed to before considering the argument. Suppose though that our reasoner doesn't respond properly to the argument and does become committed to the premises or conclusion of the argument even though she wasn't before. This is an odd thing for her to do, epistemologically speaking, in response to an argument, but it would make it possible that the reasoner would come to a decision about whether to donate in response to the argument. If the reasoner does become committed about whether to donate by accepting either a premise or the conclusion of the argument, it must be because one of the two world-norm pairs $\langle w_1, n_1 \rangle$ and $\langle w_2, n_2 \rangle$ is removed from M . Assuming that she doesn't learn anything strictly stronger than the premises or the conclusion, this requires that the part of the argument she comes to accept is false at either w_1 or w_2 . But since the parts of the argument are all silent about what ought to be the case at w_1 and w_2 (since they include those worlds paired with every possible norm), the reasoner still doesn't learn anything normative from the argument. This means that even if the reasoner acts in what appears to be an epistemically irresponsible way in response to the argument, either by accepting a premise or the conclusion without already having been committed to it, the information that the reasoner gains from the argument is non-normative. So, if the reasoner does become committed about whether to donate by accepting parts of the argument, she does so only because she is already committed to a conditional of the form 'If $w_{1(2)}$ is not actual, then I ought (not) donate.' It is this conditional, which she must already be committed to, that must play the role of the normative bridge principle in her reasoning.

the possibilities compatible with the premises, the normativity disappears. So, WORLD-NORM GAP is correct to effectively limit the scope of the is-ought gap to arguments from premises to a conjunction of the premises with the conclusion. Any normativity present in weaker conclusions is deficient for reasoning about how things ought to be for the worlds described by the premises, so it can be ignored for the purposes of understanding the gap claim.

What the above discussion brings out is that the correctness of WORLD-NORM GAP is dependent on the fact that for deductively valid arguments, the conclusion is at most as strong as the conjunction of the premises.²⁶ This property allows us to show formally that WORLD-NORM GAP is guaranteed by Gibbard's semantics. Proof: For some valid argument to C from $\{P_i\}$, suppose that each of $\{P_i\}$ is norm-invariant, $P_1 \wedge P_2 \wedge \dots$ is satisfiable, and $\{P_i\} \vdash C$. Then since the argument is valid, C must be satisfied by every world-norm pair that satisfies $P_1 \wedge P_2 \wedge \dots$. So, $P_1 \wedge P_2 \wedge \dots \wedge C$ is satisfied by every world-norm pair that satisfies $P_1 \wedge P_2 \wedge \dots$. Without loss of generality, let $\langle w, n \rangle$ be some world-norm pair that satisfies $P_1 \wedge P_2 \wedge \dots \wedge C$. Since $\langle w, n \rangle$ must satisfy each of the $\{P_i\}$ and each of the $\{P_i\}$ is norm-invariant, for each $\{P_i\} \forall n \in N(\langle w, n \rangle \in P_i)$. So $\forall n \in N(\langle w, n \rangle \in P_1 \wedge P_2 \wedge \dots)$. Then since every world-norm pair that satisfies the conjunction of the premises must satisfy the conclusion, $\forall n \in N(\langle w, n \rangle \in P_1 \wedge P_2 \wedge \dots \wedge C)$. So $P_1 \wedge P_2 \wedge \dots \wedge C$ is norm-invariant.

This proof can easily be translated into a proof of IS-UGHT GAP in any semantics of normative terms that accepts the minimal assumption that we evaluate normative sentences/utterances in terms of an ordinary possible world and normative standard. This is done by treating the possible world coordinate of evaluation like Gibbard's world, treating the normative standard like Gibbard's plan for action, and treating a sentence as non-normative iff it is norm-invariant. This last aspect of the translation should be non-controversial: if a sentence being true doesn't exclude some normative standard (which we can also think of as a normative "possibility"), then it can't put any constraints on what ought to be the case.²⁷

²⁶One might worry therefore that the understanding of the gap claim that I provide is too restricted in that it applies only to valid, non-ampliative arguments. One might think there is a plausible version of the gap claim that is meant to apply to ampliative arguments. But the most general version of this idea is clearly implausible, since we could have an ampliative logic that permits arguments from 'snow is white' to 'Rachel ought to donate to charity'. In the case of particular ampliative logics, like logics of enumerative induction, something like the is-ought gap may hold. A full understanding of those kinds of gap claims would have to proceed on an individual basis for each logic though, which is not a task I will take up here.

²⁷Note that this does require rejecting the simple view that a sentence is normative if it contains normative terms, like 'ought', 'should', or 'reason'. That simple view is certainly overly simplistic though. 'Shoplifting is permissible, or it is not', for example, should not count as normative. That said, one could accept the simplistic

What this shows is that IS-UGHT GAP should be seen as a theorem of any plausible semantics of normative language. As I argued above, IS-UGHT GAP also fits our intuitive understanding of the gap as described by Hume while taking into account the upshot of Prior's examples that the normativity must be relevant to the possibilities being reasoned about. As such, we ought to think that IS-UGHT GAP both captures and vindicates Hume's idea.

Comparing Russell and Restall's Approach

Gillian Russell and Greg Restall give a different reformulation of the is-ought gap that avoids Prior's counterexamples.²⁸ They take the gap claim to be an instance of a more general notion of an implication barrier. In their terminology, there is an implication barrier from one class of sentences to another when no sentence of the second type is implied by a collection of sentences of the first type. Restall and Russell provide a general barrier theorem that they use to derive an implication barrier from descriptive sentences to normative sentences, under a particular understanding of these classes of sentences in their models of deontic logic. The models they use are standard models consisting of a set of worlds W and a relation S supplemented with a distinguished actual world where sentences are evaluated. The relation is one of deontic accessibility, so that if all of the worlds accessible from the actual world are P -worlds, then P is obligatory. They also assume that S is transitive, euclidean, serial and secondarily reflexive, though not all of these assumptions are necessary for their proof.

Their implication barrier theorem works in cases where there is a set of sentences that are not preserved under extensions of the models or under changes of the worlds related by S . They take normative claims to be such a class. For example, while it's impermissible for Alice to hit Bob, it's permissible for Alice to hit Bob when they are in a boxing class. More formally, Russell and Restall take it that a claim is *normative* when its truth-value is always changeable either via model extension or tinkering with S in every deontic model. A claim is *descriptive* when it is always preserved under tinkering with S . It then follows as an instance of their Barrier Construction Theorem that no satisfiable collection of descriptive sentences entails a normative one.²⁹

view and still accept almost everything I say here. Their formulation of the is-ought gap would have to be in terms of non-trivial normative claims though.

²⁸ For the full details of their proposal, see Russell, Gillian, and Greg Restall. 2010. "Barriers to Implication." In *Hume on Is and Ought*, edited by Charles Pigden. Palgrave MacMillan.

²⁹ For a more precise statement of their result, see Russell and Restall, op. cit. p. 252–7.

Both IS-UGHT GAP and the formulation of the gap from Russell and Restall (2010) aim to make sense of the is-ought gap in terms of the semantics of normative language. There are many important differences between our approaches though.

Russell and Restall's approach employs Kripke models of deontic logic and claims that the class of normative sentences are those that have a certain feature with respect to that class. They treat the class of descriptive sentences similarly. Their stipulated account of normative sentences doesn't seem to capture all of the sentences that we might pretheoretically count as normative though. Peter Vranas, for example, shows that the sentence 'All citizens ought to vote' does not count as normative on their account.³⁰ Certainly though, as Vranas continues, a hypothetical argument from descriptive premises to that conclusion is in the intended domain of the is-ought gap. This appears to be a quite problematic result for taking Russell and Restall's understanding of the is-ought gap to be fully general.

The reason that 'All citizens ought to vote' doesn't count as normative on their account is because the sentence is captured in their semantics as a universally generalized disjunction. One of the disjuncts of each disjunction is non-normative, namely that the object is a citizen. Vranas shows that because the disjunction is not always unstable under both model extension and tinkering with the accessibility relation, it violates their definition of 'normative.' That said, their account also doesn't count the sentence as descriptive, as it is not always stable under model extension either. This shows that their account of the gap relies on an essentially incomplete notion of normativity — it does not count every sentence as either being normative or in the class of sentences that shouldn't entail the normative.

My account is able to count 'All citizens ought to vote' as normative, since it is not norm-invariant: that all citizens ought to vote is incompatible with a world-norm pair where the norm calls for some citizens not to vote. So, the formulation of the gap I provide would deny the validity of non-trivial arguments to this conclusion. Moreover, my account will treat any other normative claim similarly: if the claim rules out any ways the combination of the world and norms could be, it will not be norm-invariant, so there will be no non-trivial arguments from norm-invariant premises to that claim.³¹

One benefit of Russell and Restall's account is that they can state the restriction given by the is-ought gap simply: They start by delineating two classes of claims,

³⁰ See Vranas, Peter. 2010. "Comments on 'Barriers to Implication'." In *Hume on Is and Ought*, edited by Charles Pigden. Palgrave MacMillan.

³¹ Notice though that there is still a class of claims that contain normative terms that my account doesn't count as normative, namely those that are trivially true. These include conceptual truths, like 'If you ought not sit, you're not permitted to sit' but also ordinary trivialities with normative content like 'Either you ought to sit, or it's not the case that you ought to sit.' This is a benefit of the account though, as certainly, the gap claim is not meant to block inferences to these kinds of claims.

what they call 'descriptive' and 'normative,' and then they say that there are no valid arguments from sentences in the first class to a conclusion in the second. My account must be slightly more complex: it puts a restriction on what arguments are possible, but that restriction is not statable as a restriction that blocks arguments from one predetermined class of sentences to another. My formulation says that there are no valid arguments from a collection of non-normative premises to a *relevantly* normative conclusion. Which normative conclusions count as relevantly normative will depend on what the premises of the argument are. So in any particular case, the class of sentences that cannot be derived, on my view, is dependent on the premises being used. So although my account preserves our standard complete account of whether a sentence is normative, it cannot capture the gap claim simply as a restriction on arguments from one set of sentences to another.

There is a general reason to think that no account of the is-ought gap could formulate it as blocking arguments from one predetermined class of sentences to another and also cast every sentence as being in exactly one of those classes. Take a generalized version of one of Prior's examples: 'P. Therefore, P or Q.' where P is in the class with "snow is white" and Q is in the other class, which contains the paradigmatically normative sentences. For the account to put every sentence into exactly one class, it must count 'P or Q' as in the normative or in the other one. If it's normative, then 'P. Therefore, P or Q.' is a counterexample to the gap claim. If it's in the other class, then we can generate a new counterexample: 'not-P. P or Q. Therefore, Q.' So any account of the gap claim must suffer either the incompleteness disadvantage that Russell and Restall's faces or the disadvantage of not treating the is-ought gap as blocking inferences from one fixed class to another.

By taking the first horn, Russell and Restall miss out on capturing the full extent of the is-ought gap by not classing some sentences as either normative or descriptive (including some that seem like they should be). My formulation is a simple and straightforward way to take the second horn. In doing so, it is able to offer a complete way to save the gap that relies only the minimal assumption about the semantics of normative language discussed above. Hume's dictum, on my account, relies on only this weak assumption, and as such, it is much more secure than if it were to rely on the more controversial semantics provided by deontic logics that Russell and Restall use.

A final benefit of my approach over theirs is that mine offers a natural explanation of what goes wrong in Prior's proposed counterexamples. On the epistemic reading, the gap claim is about whether we can properly reason from non-normative premises to a normative conclusion. According to IS-UGHT GAP, the normativity in the conclusion of an argument is only a challenge to the is-ought gap when the normative aspects of the conclusion are relevant to the possibilities being reasoned about. So, the is-ought gap is not challenged by the purported counterexamples because no reasoner could properly employ those arguments to reach a conclusion about what ought to be the case in the situation they take themselves to be in.

Putting the gap claim in terms of a barrier to implication, as Russell and Restall do, does not offer an analogous explanation.

Conclusion

Hume's is-ought gap is a staple of introductory philosophy and logic classes, where it is typically presented as SIMPLE GAP. SIMPLE GAP is doomed by Prior's counterexamples though. Others have tried to save the claim by syntactically restricting its intended domain, like Gibbard who proposes that we only view it as a limit on arguments to simple normative conclusions. But this move is both too strong and too weak. Instead, we should seek a semantically-motivated understanding of the gap claim.

Russell and Restall's deontic-logic-based approach offers one way to do that, but as Vranas and I show, it is incomplete and its conception of normativity is overly complicated. I've introduced a simple way to understand and defend the is-ought gap that saves Hume's reasons for thinking there is such a gap. My reformulation vindicates the gap claim and offers an explanation of what goes wrong in Prior's cases. On my account, whether the is-ought gap ought to block a conclusion is relative to the particular premises in play: the is-ought gap says that we can't get from non-normative premises to conclusions about how things ought to be in cases where the premises are true.