

Barriers to Implication

Greg Restall

Philosophy Department
The University of Melbourne
restall@unimelb.edu.au

Gillian Russell

Philosophy Department
Washington University in St Louis
grussell@princeton.edu

1 The Issue

Implication barrier theses deny that one can derive sentences of one type from sentences of another. Hume's Law is an implication barrier thesis; it denies that one can derive an 'ought' (a normative sentence) from an 'is' (a descriptive sentence). Though Hume's Law is controversial, some barrier theses are philosophical platitudes; in his *Lectures on Logical Atomism*, Bertrand Russell claims:

You can never arrive at a general proposition by inference particular propositions alone. You will always have to have at least one general proposition in your premises. (Russell, 1918, p. 206)

We will refer to this claim—that one cannot derive general sentences from particular sentences—as Russell's Law.¹

A third barrier thesis claims that one cannot derive sentences about the future from sentences about the past or present. Hume's endorsement of this barrier thesis is well-known:

all inferences from experience suppose, as their foundation, that the future will resemble the past . . . if there be any suspicion that the course of nature may change, and that the past may be no rule for the future, all experience becomes useless, and can give rise to no inference or conclusion. It is impossible, therefore, that any argument from experience can prove this resemblance of the past to the future; since all these arguments are founded on the supposition of that resemblance. (Hume, EHU 4.21/37)

We will refer to this barrier thesis as Hume's Second Law. A fourth barrier thesis says that one cannot derive a necessary sentence from one about the actual world and we will refer to this last thesis Kant's Law. Such implication barrier theses present a problem.

2 The Problem

Each barrier thesis is plausible. However, each thesis leaves much unsettled. This is unsurprising, because the theses are slogans and not fully developed theories. When we attempt to settle what remains unsettled, it is no longer clear which way we can turn. There appear to be counterexamples to each of these theses, ranging from the trivial to the difficult. In this section, we consider putative counterexamples to Russell's Law because these can be found in the language of predicate logic. In this language there are good accounts of exactly what logical consequence amounts to. The same cannot be said for implication essentially involving temporal, modal or moral notions. However, we will show that the insights we gain from understanding the problem with counterexamples to Russell's Law may be applied in turn to each of the other laws. So, to putative counterexamples to Russell's Law: We all agree that the entailment from the particular Fa to the general $(\forall x)Fx$ is invalid, but what about in the special case where F is tautologous?

$$Ga \vee \neg Ga \vdash (\forall x)(Gx \vee \neg Gx)$$

Our classical tradition in logic would treat this argument as valid, because the conclusion is indeed true in every interpretation in which the premises are true—the conclusion is true in all interpretations. Of course this example seems suspicious. Tautologies are special cases. Contradictions are special cases too, and similar suspicious results obtain here:

$$Ga \wedge \neg Ga \vdash (\forall x)Gx$$

This is valid, according to the classical tradition at least, because the premise is true in no interpretations at all, so there can be no counterexamples where the premise is true and the conclusion is not. Yet the conclusion is perfectly general (it seems) and the premise is perfectly particular.

The same effect may be had without appeal to tautology or contradiction. Take this sequent, for example:

$$Fa \vdash (\forall x)(Gx \supset (\exists y)(\exists z)(Fy \wedge Gz))$$

The conclusion here seems to be universal (it states that all G s are such that something is an F and something is a G) and it certainly seems to follow (in some reasonable sense) from the premise that the object a has property F .

At the very least we must agree that stating Russell's Law precisely will involve more than a quick glance at the surface form of sentences used to express premises and conclusions. It is all too easy to infer universally quantified conclusions from premises without any quantifiers at all.²

One last example will suffice to show that our quartet of laws faces great difficulty if we seek to divide statements into two mutually exclusive and

exhaustive classes—the particular and the universal in the case of Russell’s Law; the moral and the nonmoral in the case of Hume’s Law, etc. Arthur Prior’s celebrated example takes two arguments

$$A \vdash A \vee B \quad A \vee B, \neg A \vdash B$$

where we agree that A is a nonmoral claim,³ and B is a moral claim. The problem is the status of $A \vee B$. It is either moral or nonmoral. It is moral claim, then the validity of $A \vdash A \vee B$ is a counterexample to Hume’s law. If, on the other hand, it is a nonmoral claim, then it seems that the second argument is a counterexample to Hume’s law. For the nonmoral $A \vee B$ and the nonmoral $\neg A$ (how could it be nonmoral to assert A and a moral claim to deny it?) jointly entail the moral B .⁴

* * *

One option for the defender of each barrier thesis is to retreat. Perhaps the laws as stated are not correct, but only restricted instances are correct.

Russell’s Law (Weakened) You cannot, in general, infer a universal claim from any collection of its instances. That is, you cannot, in general, infer $(\forall x)Fx$ from any set of premises only containing only instances such as Fa , Fb , Fc .

Hume’s Law (Weakened) You cannot, in general, infer a claim of the form p is obligatory from p alone.

Hume’s Second Law (Weakened) You cannot, in general, infer a claim of the form it is always the case that p from the claim that it has always been the case that p .

Kant’s Law (Weakened) You cannot, in general, infer $\Box p$ from p .

Each of these weakened forms of our barrier theses are correct but they do not come anywhere near the full power of the intended barrier theses. When Russell stated that a universal did not follow from purely particular premises, the force of the claim was not just that universal claims cannot not be inferred from any collection of purely particular claims—not just that it could not be inferred from its own instances. The case is the same for Hume’s Law: nearly no-one accepts the move from p to p is obligatory. However Hume’s Law is more controversial.

We may conclude this section: It is very difficult to state barrier theses precisely in such a way as to preserve the insights underlying each thesis, way which preserves general the force of each law. Yet, we think that something may be salvaged. To see this, however, we need to look beyond the surface

structure of each claim to what is said by that claim. Only then will we be able to both state the laws in a fashion which is both precise and justifiable, on the one hand; and with the intuitive force of the ideas motivating them, on the other.

3 The Idea

Take Russell's Law, to start. Why would one ever think it were true? Try as we might, we cannot find any justification for the Law which deals merely with the surface syntax of the claims in question. We are not led to Russell's Law from the idea that sentences involving universal quantifiers are not deducible from sentences without universal quantifiers. No-one takes it that the move from $\neg\exists x\neg Fx$ to $\forall xFx$ is a deduction from a particular to a universal. We take it that this means that $\neg\exists x\neg Fx$ is really a universal claim, despite its failure to include a universal quantifier. No, we think that the motivation for Russell's Law is properly semantic.

Thinking this way, we turn to understanding the semantics of the language of predicate logic for a motivation for Russell's Law. It is not difficult to find: consider a collection of purely premises. If each of these sentences is satisfied in some circumstance (or model) then they are also satisfied in all extensions of that circumstance (or model) which feature extra objects. If it is true in some circumstance that Fred is a frog, then if that circumstance is expanded to involve other things (frogs, non-frogs, whatever else you please) the new circumstance is also one in which Fred is a frog. Particular claims are preserved under extensions of models to include additional objects, irrespective of the properties these additional objects have. On the other hand, properly universal claims are not preserved under extensions in this way. If we have a model in which it is true that all frogs are green, we can find an extension of this model (add a new object: a non-green frog) in which it is no longer the case that all frogs are green.

This fact means that we can construct counterexamples to arguments with particular premises and a universal conclusion. Take a model satisfying all of the premises. The conclusion might be true in this model. (If it isn't, we have a counterexample to the argument already.) Since the conclusion is universal, we construct a new model, including all of the old objects from the original model, with all their original properties, but also including extra objects, such that the conclusion is now false in this new model. Since the premises were particular, they are still true in the new model, as it extended the old one. The conclusion is not true, in this model, so the argument is invalid.

Putting this intuition another way, the particular premises fix certain aspects of any model in which they are true, but they do not fix everything. The variation which remains is enough to provide counterexamples to the

properly universal claims.

We think that Russell’s Law and the notions of preservation and fragility under extension might help to explain why the other laws hold. Why can you not infer a claim that something is necessary from a collection of sentences about the actual world? Perhaps this is because necessary sentences make claims about how things are with all of the possible worlds. Why is it wrong to infer normative sentences from purely descriptive ones? Perhaps because normative sentences depend on a kind of structure in a model for normative logic, in a way that descriptive sentences do not.

We will show that these sketch semantic intuitions may be used to provide formal versions of each barrier thesis which are precise, and provable, and which also preserve the intuitions behind each thesis.

4 The Proof

We will state the theorem quite abstractly so it may apply in each of the barrier theses we have considered, and also potential cases we have not. The central formal notion will be that of a model. For the purposes of our proof, models will have two important properties. They define a notion of satisfaction, and this notion of satisfaction in turn is used to define a notion of entailment.

Definition 1 (Satisfaction) Given a formal language \mathcal{L} , for each formula A in \mathcal{L} , the model \mathfrak{M} will either satisfy that formula (written ‘ $\mathfrak{M} \models A$ ’) or it will not satisfy that formula (written ‘ $\mathfrak{M} \not\models A$ ’).

Examples of models in this sense are Tarski’s models for the language of predicate logic, and possible worlds models for the language of modal logic. Notice, however, that given a possible worlds model in which formulas take differing truth values in different worlds, we must make a choice of what counts as the entire model satisfying a formula. When it comes to this, we will choose a simple technique. We will pick one possible world in our model to count as the “actual world,” in the sense that the formulas true in that world are those said to be satisfied by the model.

Models arise in the analysis of logical consequence. An argument is valid, or the conclusion is said to be entailed by the premises just when every model satisfying the premises also satisfies the conclusion:

Definition 2 (Entailment) Given a class \mathcal{M} of models, we will say that a collection X of premises \mathcal{M} -entails a conclusion A just when every model \mathfrak{M} in \mathcal{M} satisfying each of the elements of X , also satisfies A . This is written as ‘ $X \vdash_{\mathcal{M}} A$ ’ or when no confusion about the class \mathcal{M} may arise, as ‘ $X \vdash A$.’

Again, if \mathcal{M} is the class of all Tarskian models of the language of predicate logic, then $\vdash_{\mathcal{M}}$ is classical logical consequence. If \mathcal{M} is the class of all possible worlds models of some kind, the consequence relation will be some corresponding kind of modal logic.

The rest of our definitions will consider a given class \mathcal{M} of models, and a relation R between models. For example, \mathcal{M} could be the class of Tarskian models for first order logic and R the relation of model extension.

Definition 3 (R-Preservation) A formula A is R -preserved if and only if

$$(\forall \mathfrak{M}, \mathfrak{M}' \in \mathcal{M})(\text{if } \mathfrak{M} \models A \text{ and } \mathfrak{M} R \mathfrak{M}' \text{ then } \mathfrak{M}' \models A)$$

In other words, if a model \mathfrak{M} satisfies A , and \mathfrak{M} is related by R to \mathfrak{M}' then \mathfrak{M}' satisfies A too.

If, as in the example above, R is the relation of model extension, then the R -preserved formulas are those which are semantically particular.

Note this curious fact: The failure of R -preservation is not necessarily a sign that we have a semantically universal formula on our hands. Take the example of $Fa \vee (\forall x)Gx$. This is sometimes R -preserved (if you have a model in which Fa is satisfied, $Fa \vee (\forall x)Gx$ is satisfied in any extension of it). However, it is sometimes not (take a model in which Fa is false, but $(\forall x)Gx$ is true—extend it to a model in which G fails of some objects).

So, what is a formal notion representing the kind of universality we seek? It is not too difficult to find: it is fragility.

Definition 4 (R-Fragility) A formula A is R -fragile if and only if

$$(\forall \mathfrak{M} \in \mathcal{M})(\text{if } \mathfrak{M} \models A \text{ then } (\exists \mathfrak{M}')(\mathfrak{M} R \mathfrak{M}' \text{ and } \mathfrak{M}' \not\models A))$$

At last, we can present our theorem.

Theorem 5 (Barrier Construction Theorem) Given a class \mathcal{M} of models, and a collection $X \cup \{A\}$ of formulas. If (a) X is satisfied by some model in \mathcal{M} ; (b) A is R -fragile; and (c) each element of X is R -preserved, then $X \not\vdash_{\mathcal{M}} A$.

We call this the “Barrier Construction Theorem” because it shows that any R relation on models sets up a barrier that implication cannot cross. Its proof is straightforward:

Proof Since X is satisfied by some model (a), choose one such model, \mathfrak{M} . If $\mathfrak{M} \not\models A$, then $X \not\vdash_{\mathcal{M}} A$ and we are finished. On the other hand, if $\mathfrak{M} \models A$, then since A is R -fragile (b), there is some \mathfrak{M}' where $\mathfrak{M} R \mathfrak{M}'$ and $\mathfrak{M}' \not\models A$. Now, since each element of X is R -preserved (c), \mathfrak{M}' satisfies each element of X , and \mathfrak{M}' is our counterexample to the validity of the argument from X to A : $X \not\vdash_{\mathcal{M}} A$. \square

This theorem has a powerful converse, which indicates that the pairing of preservation and fragility precisely capture inferential barriers.

Definition 6 (barrier) $\langle \Sigma | \Delta \rangle$ is a barrier if and only if whenever $\Gamma \subseteq \Sigma$ is a satisfiable subset of Σ and $B \in \Delta$, then $\Gamma \not\vdash B$

Barriers are not necessarily partitions of a language, but they do demarcate a gulf unbridgable by valid implication. We have shown that relations among models can be used to define barriers. We now show that, if a consequence relation is defined by truth preservation on a class of models, then every barrier with respect to that consequence relation is generated by some relation R in this way.

Theorem 7 (Barrier Representation Theorem) Given a barrier $\langle \Sigma | \Delta \rangle$, there is some binary relation R on models where each element of Σ is R -preserved, and each element of Δ is R -fragile. Furthermore, R can, without loss of generality, be presumed to be reflexive and transitive.

Proof Define the binary relation R on models as follows: $\mathfrak{M} R \mathfrak{M}'$ if and only if $(\forall A \in \Sigma)(\mathfrak{M} \models A \Rightarrow \mathfrak{M}' \models A)$. Then it first follows that each element of Σ is R -preserved. To show that each element in Δ is R -fragile, let $B \in \Delta$. If $\mathfrak{M} \models B$, then let $\Gamma = \{A \in \Sigma : \mathfrak{M} \models A\}$. $\Gamma \not\vdash B$ (as $\langle \Sigma | \Delta \rangle$ is a barrier), so let \mathfrak{M}' be a model in which Γ is true but B is not. $\mathfrak{M} R \mathfrak{M}'$, so B is R -fragile.

The relation R we have defined is reflexive ($\mathfrak{M} R \mathfrak{M}$, for each \mathfrak{M}) and transitive (if $\mathfrak{M} R \mathfrak{M}'$ and $\mathfrak{M}' R \mathfrak{M}''$ then $\mathfrak{M} R \mathfrak{M}''$, for each \mathfrak{M} , \mathfrak{M}' and \mathfrak{M}''). So, every barrier can generated by a reflexive and transitive relation. \square

5 The Applications

With those theorems at hand, we will turn to our barrier theses. We will cast them as claims of the existence of a barrier in the sense of Definition 6. Given the Barrier Representation Theorem, we will look for a binary relation on models which can be used to define the barrier which is our target, in each case.

5.1 Generality

We begin with the most familiar case: Russell's Law. Our goal here is to distinguish particular from general sentences in such a way that the distinction respects our intuitions about which sentences are semantically general or particular. It will follow that a simple version of Russell's law comes as not only true, but provable.

Our class of models (\mathcal{U}) will be the class of Tarskian models for first order logic. Let the relation R between models be the relation of model extension (\supseteq), where a model \mathfrak{M}' extends a model \mathfrak{M} just in case \mathfrak{M}' can

be obtained from \mathfrak{M} by adding more objects to the domain and extending the interpretation of the predicates to cover the cases of the new objects. (Formally speaking, if F an n -place predicate and α an assignment of variables to values in the domain of \mathfrak{M} (avoiding the extra objects in \mathfrak{M}') then $\mathfrak{M}, \alpha \models Fx_1, \dots, x_n$ if and only if $\mathfrak{M}', \alpha \models Fx_1, \dots, x_n$.)

Definition 8 (semantic particularity) A sentence is semantically particular iff it is \supseteq -preserved, that is, for each $\mathfrak{M}, \mathfrak{M}' \in \mathcal{U}$, if $\mathfrak{M} \models A$ and $\mathfrak{M}' \supseteq \mathfrak{M}$ then $\mathfrak{M}' \models A$.

Definition 9 (semantic universality) A sentence is semantically universal iff it is \supseteq -fragile, that is, for each $\mathfrak{M} \in \mathcal{U}$ where $\mathfrak{M} \models A$, there is some $\mathfrak{M}' \in \mathcal{U}$ such that $\mathfrak{M}' \supseteq \mathfrak{M}$ and $\mathfrak{M}' \not\models A$.

Semantically universal sentences include $(\forall x)Fx$, $(\forall x)Fx \wedge p$ and $(\forall x)(Fx \supset Gx)$. Semantically particular sentences include Fa , $\neg Fa$, $Fa \wedge Ga$, $(\exists x)Fx$ and logical truths. The place of this later kind of sentence in the category is easier to understand when we recall that we are thinking of semantically particular sentences as those which fail to constrain the whole world in the way that universal sentences do.

Some sentences are neither semantically universal nor semantically particular. $Fa \vee (\forall x)Gx$ is not semantically particular because given a model \mathfrak{M} such that $\mathfrak{M} \models (\forall x)Gx$ (and hence $\mathfrak{M} \models Fa \vee (\forall x)Gx$) but $\mathfrak{M} \not\models Fa$ we can find a model \mathfrak{M}' such that $\mathfrak{M}' \supseteq \mathfrak{M}$ and $\mathfrak{M}' \not\models Fa \vee (\forall x)Gx$. Moreover the disjunction is not semantically universal because where $\mathfrak{M} \models Fa$, $\mathfrak{M} \models Fa \vee (\forall x)Gx$, but all extensions of \mathfrak{M} will be models which satisfy the disjunction since they will satisfy Fa .

In addition, another boundary condition for these definitions are the sentences are both semantically universal and semantically particular. These are sentences such that, if true in a model, they are true in all extensions of that model, and furthermore, there is some extension in which they are not true. It follows that these are sentences which are true in no model at all. These are the inconsistent sentences.

With the definition of the model-extension relation and the barrier defined in terms of this relation, we are in a position to state and prove Russell's Law:

Corollary 10 (Russell's Law) If Σ is a satisfiable set of sentences, each of which is semantically particular, and A is semantically universal, then $\Sigma \not\models A$.

Proof Russell's Law is an instance of the Barrier Construction Theorem. \square

The putative counterexamples to the law are defused by attention to the type (universal, particular or neither) of the sentences involved. Prior's argument, $Fa \vdash Fa \vee (\forall x)Gx$, does not have a semantically universal conclusion, and nor does $Fa \vdash (\forall x)(Gx \supset (\exists y)(\exists z)(Fy \wedge Gz))$: consider a model in

which neither Fx nor Gx are true of any object and extensions of that model in which Gx is true of some objects but Fx is true of none.

5.2 Necessity

Now we turn to the modal case, taking propositional $s5$ as our example both for its simplicity and its strength. Our models in this case are sets of worlds, on which is defined a reflexive, symmetric and transitive accessibility relation $\langle \mathcal{W}, S, g \rangle$. These models contain a privileged world g , (the actual world) used for defining truth in the model:

$$\mathfrak{M} \models A \quad \text{if and only if} \quad g \Vdash A$$

This time the relation R will be that of modal model extension (\sqsupseteq). A model \mathfrak{M}' is an extension of a model \mathfrak{M} in this sense if it can be obtained from \mathfrak{M} by adding new worlds and extending the accessibility relation (we may add new pairs of worlds to R , but we are not allowed to take any away). Now we use this notion of modal model extension to define two classes of sentence:

Definition 11 (modal particularity) A sentence A is modally particular iff it is \sqsupseteq -preserved, that is, for each $\mathfrak{M}, \mathfrak{M}' \in \mathcal{V}$, if $\mathfrak{M} \models A$ and $\mathfrak{M}' \sqsupseteq \mathfrak{M}$ then $\mathfrak{M}' \models A$.

Definition 12 (modal generality) A sentence A is modally general iff is \sqsupseteq -fragile, that is, for each $\mathfrak{M} \in \mathcal{V}$ where $\mathfrak{M} \models A$, there is some $\mathfrak{M}' \in \mathcal{V}$, such that $\mathfrak{M}' \sqsupseteq \mathfrak{M}$ and $\mathfrak{M}' \not\models A$.

The class of modally particular sentences includes sentences from three groups: (i) the sentences which contain no modal operators, such as p , and $p \vee q$, (ii) sentences such as $\diamond p$ and $\neg \Box p$ whose truth can be secured by a single world, or a structure of worlds, regardless of additions to that structure. Since we are dealing with $s5$, $\Box \diamond p$ falls into this category. And (iii) the logical truths of modal logic, such as $\Box p \vee \neg \Box p$.

The modally general sentences include $\Box p$, $\Box(p \vee q)$ and $\neg \diamond p$. The disjunction of a modally particular sentence with a modally general one is neither modally general nor particular, and unsatisfiable sentences are both modally general and modally particular.⁵

Corollary 13 (Kant's Law) If Σ is a satisfiable set of sentences, each of which is modally particular, and A is modally general, then $\Sigma \not\models A$.

Proof This is an instance of the Barrier Construction Theorem. □

5.3 Time

In the case of Hume’s Second Law our models will be sets of time slices (w_i) ordered by a relation of temporal precedence $\langle \mathcal{W}, <, w_p \rangle$ with a distinguished moment (w_p), (which can be thought of informally as the present moment), used for defining truth in the model:

$$\mathfrak{M} \models A \quad \text{if and only if} \quad w_p \Vdash A$$

We let the R relation between models be the symmetric relation of history sharing (Υ) where two models stand in this relation if they are the same with respect to the present moment and all earlier moments.

Definition 14 (semantic historical sentences) A sentence A is semantically historic iff it is Υ -preserved, that is for each $\mathfrak{M}, \mathfrak{M}' \in \mathcal{T}$, if $\mathfrak{M} \models A$ and $\mathfrak{M}' \Upsilon \mathfrak{M}$ then $\mathfrak{M}' \models A$.

Definition 15 (semantically future-constraining) A sentence A is semantically future-constraining iff it is Υ -fragile, that is, for each $\mathfrak{M} \in \mathcal{T}$ where $\mathfrak{M} \models A$, there is some $\mathfrak{M}' \in \mathcal{T}$, such that $\mathfrak{M}' \Upsilon \mathfrak{M}$ and $\mathfrak{M}' \not\models A$.

Suppose that the $<$ relation is transitive, irreflexive and anti-symmetric. Then p , Pp , Hp , GPp ⁶ are semantically historic, Fp and Gp are semantically future-constraining and both $p \vee Fp$ and PFp are neither semantically historic nor semantically future-constraining.

Now we can formulate Hume’s second law on the model of the previous barrier theses:

Corollary 16 (Hume’s Second Law) If Σ is a satisfiable set of sentences, each of which is semantically historic, and A is semantically future-constraining, then $\Sigma \not\models A$.

Proof Once again the law follows from theorem 5. □

5.4 Normativity I

The original Hume’s Law is our most interesting case—an barrier thesis that some philosophers take to be false. Our task, as before, is to choose the right inter-model relation for defining ‘ought’ sentences and prior to that, we will need a kind of model relevant to evaluating the inferential properties of normative sentences. We will use models for a straightforward deontic logic: sets of worlds on which a binary relation of moral satisfaction is defined, and containing a distinguished actual world for defining truth in the model $\langle \mathcal{W}, S, g \rangle$. Call the class of these models ‘ \mathcal{N} ’. Semantics for an ought operator (O) are given by the clause:

$$w \Vdash OA \quad \text{if, and only if, for all } w' \text{ such that } wSw', w' \Vdash A$$

And for a permissibility operator (P) by

$w \Vdash PA$ if, and only if, for some w' such that $wSw', w' \Vdash A$

By assuming different properties of the relation of moral satisfaction between worlds, we obtain stronger and weaker deontic logics. We will assume that S is transitive, euclidean, serial and secondarily reflexive, which makes the following arguments truth-preserving:

$OA \vdash OOA$

$\neg OA \vdash O\neg OA$

$OA \vdash \neg O\neg A$

$\vdash O(OA \supset A)$

It will be interesting to see whether we can formulate Hume's Law given such a strong deontic logic.

In trying to divide up the sentences of the language of deontic logic we run into a question about the interpretation of Hume's Law. Does it say that one cannot get 'ought'-sentences from 'is'-sentences? Or does it say that one cannot get normative sentences from descriptive sentences? Sentences that ascribe permissibility, such as Pp are normative sentences, but they are not 'ought'-sentences, so there is a question about which side of the divide we want to put them on. Hume himself was worried about how to derive sentences containing the words 'ought' and 'ought not' from those containing the words 'is' and 'is not', as the following quote shows:

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 way 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 copulation 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. (T 3.1.1.27/469–70)

Nonetheless it is very common to hear the law glossed as saying that one cannot get normative sentences from descriptive ones. We think both interpretations are defensible, and we will formulate both.

Definition 17 (normative extension: \ni) $\mathfrak{M}' \ni \mathfrak{M}$ just in case you can obtain \mathfrak{M}' from \mathfrak{M} by adding new worlds and extending the S relation.

Definition 18 (normative translation: $\check{\ni}$) $\mathfrak{M}' \check{\ni} \mathfrak{M}$ just in case you can obtain \mathfrak{M}' from \mathfrak{M} simply by changing the pairs of worlds related by the S relation.

Definition 19 (normative particularity) A sentence A is normatively particular iff it is \ni -preserved, that is, for each $\mathfrak{M}, \mathfrak{M}' \in \mathcal{N}$, if $\mathfrak{M} \models A$ and $\mathfrak{M}' \ni \mathfrak{M}$ then $\mathfrak{M}' \models A$.

Definition 20 (normative generality) A sentence A is normatively general iff it is \ni -fragile, that is, for each $\mathfrak{M} \in \mathcal{N}$ where $\mathfrak{M} \models A$, there is some $\mathfrak{M}' \in \mathcal{N}$ such that $\mathfrak{M}' \ni \mathfrak{M}$ and $\mathfrak{M}' \not\models A$.

Definition 21 (descriptiveness) A sentence A is descriptive iff it is $\check{\ni}$ -preserved, that is, for each $\mathfrak{M}, \mathfrak{M}' \in \mathcal{N}$, if $\mathfrak{M} \models A$ and $\mathfrak{M}' \check{\ni} \mathfrak{M}$ then $\mathfrak{M}' \models A$.

Definition 22 (normativity (sufficient condition)) A sentence A is normative iff it is $\check{\ni}$ -fragile, that is, for each $\mathfrak{M} \in \mathcal{N}$ where $\mathfrak{M} \models A$, there is some $\mathfrak{M}' \in \mathcal{N}$ such that $\mathfrak{M}' \check{\ni} \mathfrak{M}$ and $\mathfrak{M}' \not\models A$.

As in the modal case, there are three kinds of sentence that turn out to be normatively particular: (i) sentences containing no deontic operators, such as p and $p \supset q$; to evaluate these we need only look at the actual world, (ii) sentences containing deontic operators whose truth can be secured by some structure of worlds at least one of which is in the relation of moral satisfaction to the actual world (regardless of what else we add to that structure), such as Pp , PPp and $\neg Op$, and (iii) the deontic tautologies—the things that are \ni -preserved only because there is no world in any model where they are false.

The normatively general sentences include Op and $\neg Pp$. Given any model which satisfies these sentences we can always extend it to make them false.

The descriptive sentences are, intuitively, those that make no appeal to the deontic structure of the model (at least, not the parts that can change. Deontic tautologies will be descriptive.) They include p , $p \vee q$ and $Op \vee \neg Op$.

Finally, the normative sentences are, intuitively, the ones which make demands on the arrangement of the relation of moral satisfaction. They include Pp . Oddly though, important normatively general sentences such as Op are not $\check{\ni}$ -fragile. Consider a model in which every world is one where p is true. No rearranging of the S relation can make it the case that Op is false, and so Op is not $\check{\ni}$ -fragile (though it is not $\check{\ni}$ -preserved either).

It seems that fragility under normative extensions captured what was special about ought-sentences; they constrained entire models, and fragility

under normative transformations captured what was special about permissibility-sentences; they require the existence of a particular kind of structure of worlds standing in the S relation to the actual one. Normative sentences then, would be those that are either normative extension fragile, or normative transformation fragile.

Definition 23 (Normativity) A sentence A is normative iff it is either \checkmark -fragile or \ni -fragile, that is, either i) for each $\mathfrak{M} \in \mathcal{N}$ where $\mathfrak{M} \models A$, there is some $\mathfrak{M}' \in \mathcal{N}$ such that $\mathfrak{M} \checkmark \mathfrak{M}'$ and $\mathfrak{M}' \not\models A$ or ii) for each $\mathfrak{M} \in \mathcal{N}$ where $\mathfrak{M} \models A$, there is some $\mathfrak{M}' \in \mathcal{N}$ such that $\mathfrak{M}' \ni \mathfrak{M}$ and $\mathfrak{M}' \not\models A$

Finally we can formulate our two versions of Hume's Law:

Corollary 24 (Hume's Law (Ought-formulation)) If Σ is a satisfiable set of sentences, each of which is normatively particular, and A is normatively general, then $\Sigma \not\models A$.

Proof This follows from the Barrier Construction Theorem. \square

Corollary 25 (Hume's Law (Normativity-formulation)) If Σ is a satisfiable set of sentences, each of which is descriptive, and A is normative, then $\Sigma \not\models A$.

The ought-formulation follows from the Barrier Construction Theorem as usual, but what about the normativity formulation? If we could show that no descriptive sentences imply normatively general ones, then the rest of law would follow by the Barrier Construction Theorem (since it is an instance of the Barrier Construction Theorem that no descriptive sentences imply sentences not preserved under normative transformation). We show that all descriptive sentences are normatively particular:

Lemma 26 All descriptive sentences are normatively particular.

Proof Suppose a sentence A is not normatively particular, that is, not \ni -preserved. Then there are models \mathfrak{M} and \mathfrak{M}' such that $\mathfrak{M}' \ni \mathfrak{M}$, $\mathfrak{M} \models A$ and $\mathfrak{M}' \not\models A$.

Let \mathfrak{M}^* be any \ni -extension of \mathfrak{M} which adds just those worlds that are in \mathfrak{M}' but does not relate any of the new worlds to any of the original worlds in \mathfrak{M} . Then, since $\mathfrak{M} \models A$, $\mathfrak{M}^* \models A$ too; the new worlds do not influence the evaluation on the original worlds. But $\mathfrak{M}^* \checkmark \mathfrak{M}'$, and $\mathfrak{M}' \not\models A$, so A is not \checkmark -preserved, and hence not descriptive.

Since A was an arbitrary sentence, it follows that any sentence that is not normatively particular is not descriptive. Contraposing, if a sentence is descriptive it is normatively particular. \square

We can use this fact to prove the normativity formulation of Hume’s Law.

Proof (a) No descriptive sentence entails a \exists -fragile one. (This is an instance of the Barrier Construction Theorem). (b) All descriptive sentences are \exists -preserved, (by Lemma 26) and no \exists -preserved sentence entails a \exists -fragile sentence (by the Barrier Construction Theorem), so no descriptive sentence entails a \exists -fragile sentence. Therefore, descriptive sentence entails a normative one (from (a) and (b), using the definition of ‘normativity’.) \square

5.5 Normativity II

Some philosophers believe that the use of deontic logics to tackle ethical and metaethical issues is a mistake. Suppose for the moment that they are right. We still think that the intuitive ideas of preservation and fragility over changes in situation explain Hume’s Law, and moreover, deserve further attention from philosophers. One reason we think this is that the idea that normative sentences are fragile over changes in situation through which descriptive sentences are preserved is quite plausible. A second is that the idea is powerful enough to have some interesting consequences. We argue for both these points in the rest of this section.

First, the idea that normative sentences are fragile over changes in situation through which descriptive sentences are preserved is plausible. Consider a very simple situation in which the only thing that happens is that Alice intentionally hits (and hurts) Bob. This is *prima facie* a situation in which Alice does something wrong. But now ‘extend’ (in an informal sense) the situation to one in which Alice and Bob are in training for a boxing tournament, and Bob demands (in the name of good training) that his partners try as hard as they can to win. Suddenly it is plausible that what Alice did was not wrong after all, though the description we gave of the first situation is still a true one. Now imagine a third situation, an extension of both the first and second ones. This time we add Candy, a suicidal anti-boxing protester who has informed Alice that she will kill all three of them if Alice hits Bob. Surely Alice would be wrong to hit Bob in that situation. And yet the description we gave of the second situation is still true of this third one. So we think it plausible that the sentence ‘it is wrong for Alice to hit Bob’ is fragile over these changes in situation, though (for example) ‘Alice hits Bob’ is not.

As it stands this is just one sequence of possible situations, each extending the last. But the story is a familiar one from the discipline of normative ethics. Someone suggests a principle to explain an ethical intuition—e.g. whenever the situation is *F*, one ought to do *A*, and a counterexample to the principle is quickly unearthed. We think that the explanation for this might be that the principle links preserved with fragile sentences.

We can make the same sort of case for the fragility of non-ethical normative sentences. Consider sentences about epistemic justification. Suppose that I am wandering around Christchurch on my first ever visit to New Zealand, and I come across a strange bird in the street. It looks like a kiwi, and, knowing that kiwis are very rare and that the only kiwis in Christchurch are in the house by the aquarium, I form the belief that a kiwi has escaped from the kiwi house next to the aquarium. This belief seems at least as justified as many of my ordinary beliefs. But now extend the former situation by supposing that it is the case that, and I believe to be the case that, the kiwi has a relation—the kiwik—which looks very similar. Similar enough that it takes some expertise in antipodean ornithology to tell the kiwi apart from the kiwik (expertise which I do not have). Moreover, unlike the kiwi, the kiwik is a common sight in urban Christchurch. Surely the justification for my belief that a kiwi has escaped from the kiwi house—my sighting of a kiwi-like bird in the street—has been undermined, though the original evidence has not been removed, but merely extended.

Secondly, we think that the idea that normative sentences are fragile with respect to extensions where descriptive sentences are not might have interesting consequences. We suggest two here:

First, if we come to believe that normative sentences are fragile this could have a role to play in deciding controversial cases of normativity. It could clarify, for example, the debate over whether metalinguistic meaning sentences, such as ‘horse means horse’ are semantically normative.

Second, the thesis may have consequences for the epistemology of the normative, for example, for the question of whether our beliefs in normative propositions can be justified through perception or intuition. The epistemology of the normative is difficult. Some philosophers have suggested that the mechanism by which our beliefs in normative propositions are justified is, or is similar to, perception. But suppose that we can make it plausible that the content of perception is limited to propositions expressible by sentences which are preserved over certain changes—changes over which normative sentences are fragile.⁷ Then how can be that the fragile normative property is perceivable? Such a thought also motivates a sceptical challenge: what reason could we have to believe that we are in a situation in which a normative sentence is true rather than a situation in which it is false if both are consistent with our situation-change preserved perceptual evidence?

This question has a Cartesian flavour, and we do not mean to suggest that it is unanswerable. We only mean to illustrate that the fragility of the normative with respect to the descriptive is not just plausible, but is also of philosophical interest—even for those who doubt the helpfulness of deontic logics in thinking about ethics.

6 The Conclusion

In conclusion we note that our view on Hume's law has consequences for a number of issues surrounding the law.

First, Charles Pigden (Pigden, 1989) has argued that if Hume's Law is true, then normative terms cannot be logical constants, as they are taken to be in the standard deontic logics; such logics license apparently Hume's Law violating arguments such as $\neg(p \wedge \neg p) \vdash O\neg(p \wedge \neg p)$. In our view taking normative terms as logical constants is consistent with accepting Hume's Law. Once the uses of 'normative' and 'descriptive' in the statement of the law are properly understood, the law is seen not to conflict with this kind of argument. In this case, $O\neg(p \wedge \neg p)$ is not really a normative sentence.

As a number of writers on Hume's Law have noted, the converse of Hume's Law does not have the same intuitive pull as the law itself, and informal arguments in the literature suggest that you can get an 'is' from an 'ought'. On our view this asymmetry has an explanation in terms of preservation and fragility under extensions. Whilst one can't derive R-fragile sentences from R-preserved sentences, the converse is clearly not true. Here are just two counterexamples: $(\forall x)Fx \vdash Fa$, (a \supseteq -preserved sentence from a \supseteq -fragile one) and $\Box P \vdash P$, (a \supseteq -preserved sentence from a \supseteq -fragile one). So on our account it is understandable that, while one can get an 'is' from an 'ought', one still cannot get a genuine 'ought' from an 'is'.⁸

Notes

¹We name the implication barrier theses after philosophers who have maintained them for ease of reference only, and make no claims about who first formulated each law.

²The same point obtains for the other Laws, but we will not stop to spell out the details here. It suffices to change universal quantifiers to claims of necessity (all possible worlds) or the future (at all future moments) or obligation (at all permissible circumstances) and the same kinds of results will obtain.

³In the case of Hume's Law, which was Prior's target.

⁴It might be thought that a move to a relevant consequence relation might avoid the trouble by denying either disjunctive addition or disjunctive syllogism. While relevant logics have much to commend them, a solution to Prior's puzzle is not amongst their virtues (Humberstone, 1996).

⁵ It is interesting to note what happens when we add identity to the logic in the modal case. Entailments involving identity sentences might be thought to provide counterexamples to Kant's Law, since in some modal logics the following sequent is necessarily truth preserving:

$$a = b \vdash \Box(a = b)$$

Is this a counterexample to Kant's Law? No, because a model which satisfies $a = b$ will only have extensions in which $a = b$ is true at every world. This makes both identity sentences and identity sentences with a ' \Box ' on the front modally particular, and hence the implication is no counterexample to

Kant's Law.

⁶In what follows we exploit Prior's tense logical operators 'G', 'H' 'F' and 'P' whose semantics is given by the clauses ' $w \Vdash GA$ iff $\forall w_i w < w_i, w_i \Vdash A$ ' ' $w \Vdash HA$ iff $\forall w_i, w_i < w, w_i \Vdash A$ ', ' $w \Vdash FA$ iff $\exists w_i, w < w_i w_i \Vdash A$ ', and ' $w \Vdash PA$ iff $\exists w_i, w_i < w, w_i \Vdash A$ ', respectively.

⁷One example might involve a situation in which you torture a robot dog that feels pain, and a situation in which you torture a robot dog that merely acts as if it feels pain. Another might involve an ordinary situation in which you consider giving money to help feed the hungry and one where you consider giving money to help feed and, unbeknownst to you, your action will set in motion a chain of events that causes even more wide-spread famine than before.

⁸Thanks to audiences at the University of Melbourne, Monash University and the Australian National University, and especially to Lloyd Humberstone for helpful comments on this paper.