Foundations Seminar: Barriers to Entailment

Gillian Russell, Semester 2, 2024

Week 4: Universality and Time

Six Tests from last week

1. Can it get past Prior's dilemma?
2. Classification of Vranas sentences?

$$p \to Oq, \quad p \vee Oq$$

3. *Ought* implies *can* (complex contraposition more generally)

$$Op \vDash \Diamond p \qquad \neg\Diamond p \vDash \neg Op$$

$$Fp \vDash \Diamond p \qquad \neg\Diamond p \vDash \neg Fp$$

4. Informal arguments (speech acts, quotation, thick normative expressions)
5. demarkation of normative/descriptive
6. unified treatment of the barriers

The Universal Barrier

— Our first task is to say what it is for sentences to be *Particular* and *Universal*.

— Central idea: true particular claims *stay true* when the model is extended:

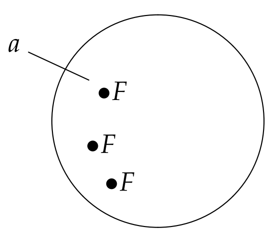

Fa is true

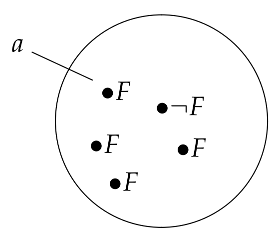

Fa is still true

Whereas true Universal claims can be made false by extending the model:

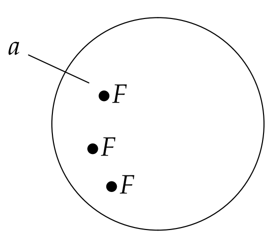

$\forall x Fx$ is true

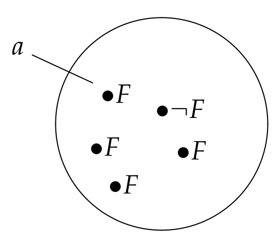

$\forall x Fx$ is no longer true

— We're using the relation of model-extension to define Particular and Universal. We can tighten this up a bit:

Definition (Model-Extension ($\Subset$)). *A model N extends a model M just in case*

1. $D_M \subseteq D_N$ *and*
2. *for all individual constants* α,

$$I_N(\alpha) = I_M(\alpha)$$

3. *for all n-place predicates* Π,

$$I_M(\Pi) \subseteq I_N(\Pi)$$

Definition (R-preserved and -Fragile). *A sentence is*

– R-preserved *iff any model which satisfies it is such that all R-related models satisfy it too.*
– R-fragile *iff all models which satisfy it are R-related to at least one model which does not.*

Definition (Particular and Universal Sentences). *A sentence is*

– *Particular iff it is* $\Subset$*-preserved.*
– *Universal iff it is* $\Subset$*-fragile.*

— Let's try to understand these definitions better by looking at some examples.

Sentences which are $\Subset$ -preserved (Particular):

$$Fa, \quad \neg Fa, \quad Fa \vee Gb, \quad Fa \to Gb, \quad \exists x Fx, \quad \neg\forall x Fx, \quad \forall x(Fx \vee \neg Fx)$$

Sentences which are $\Subset$-fragile (Universal) :

$$\forall x Fx, \quad \neg\exists x \neg Fx, \quad \forall x \forall y Rxy, \quad \forall x(Fx \wedge Gx)$$

Sentences which are Neither:

$$Fa \vee \forall x Gx, \quad \neg Fa \to \forall x Gx$$

Sentences which are Both(!):

$$Fa \wedge \neg Fa$$

Then the theorem says:

(R&R) No satisfiable set of R-preserved sentences entails an R-fragile sentence

This turns out to be easy to prove!

Proof 1. *Suppose the set of premises Γ is satisfiable and the conclusion ϕ is R-preserved. Then Γ has a model. Does that model make ϕ true? If not, we have a counterexample. If it does, that model has an R-related model which makes the ϕ false (from the definition of R-fragile) and that very model satisfies Γ (from the definition of R-preserved). So either way there is a model that makes Γ true and ϕ false, so $\Gamma \not\models \phi$.*

— What will this say about the Universal analogue of Priori (1)?

$$Fa \models Fa \vee \forall x Gx$$

— What about Prior (2)?

$$Fa \vee \forall x Gx, \neg Fa \models \forall x Gx$$

Test	R&R
Prior's Dilemma	✓
Classification of Vranas sentences	✗
Ought implies *Can*	✗
informal (Searle, thick n.e., quotation etc.)	?
demarkation of normative/descriptive	✗
unified treatment of barriers?	?

Particular and Universal sets

— Above we saw sentences that are Neither:

$$Fa \vee \forall x Gx$$

$$Fa \rightarrow \forall x Gx$$

— Peter Vranas has argued that this classification is problematic: and makes the R&R barrier too weak. This is because the mixed Vranas conditionals ($Fa \rightarrow \forall x Gx$), when taken together with their Particular antecedents (Fa), entail their Universal consequents ($\forall x Gx$.) This strongly suggests that the conditionals themselves should be classified as Universal if we are to capture the full strength of the barrier.

— But we *can't* classify them as Universal. If we did, then the barrier thesis would be false, since

$$\neg Fa \models Fa \rightarrow \forall x Gx$$

and $\neg Fa$ is Particular. This is Vranas' problem.[1]

The point is especially clear in the normative case. *If Fred is five years old, then Fred ought not to have homework* seems unassailably normative, and especially obviously so when we consider adding the descriptive antecedent to the premise set, since the resulting set then entails *Fred ought not to have homework.*

[1] Vranas, P. (2010). Comment on 'Barriers to Implication'. In Pigden, C., editor, *Hume, Is and Ought: New Essays*, pages 260–267. Palgrave MacMillan

— It almost seems tempting to say that the status of a mixed conditional varies with the company it keeps. When it hangs around with Fa it "acts" universal. But when it spends time with $\neg Fa$ it "acts" particular. Similarly for mixed disjunctions.

— Some philosophers have noticed this while thinking about the *is/ought* barrier and suggested that whether a sentence is normative might depend on context.

— I want to suggest an alternative approach

— It builds on an idea from Philippa Foot that we first encountered a couple of weeks ago.

> We cannot possibly say that at least one of the premises must be evaluative if the conclusion is to be so; for there is nothing to tell us that whatever can truly be said of the conclusion of a deductive argument can truly be said of any one of the premises. It is not necessary that the evaluative element should "come in whole" so to speak. If f has to belong to the premises it can only be necessary that it should belong to the premises *together*, and it may be no easy matter to see whether a set of propositions has the property f. [2] (507)

— reinforced by informal "quotation" counterexamples

— The first step is to note is that model-theoretic properties can sometimes be extended quite naturally from sentences to sets of sentences.

— We have already done this with with the property of truth-in-a-model: a set of FOL-sentences is true-in-a-model just in case each of the set's members is true-in-that-model.

— But sometimes a set may have a property even if not every member of the set has it, and sometimes even if even if none of the members has the property individually.

— A sentence is *unsatisfiable* if there is no model which makes it true, and a set of sentences unsatisfiable if there is no model which makes *it* true (by making all of its members true.) The set $\{Fa, \neg Fa\}$ is unsatisfiable, but none of its members is.

— In this spirit, we will extend the definitions of Particularity and Universality in terms of fragility and anti-fragility from sentences to sets:

Definition (Particularity (sets))**.** *A set of sentences, Γ, is P if and only if it is $\Subset$–anti-fragile, i.e.iff, if $V_M(\Gamma) = 1$, then for all extensions of M, N, $V_N(\Gamma) = 1$ too.*

Definition (Universality (sets))**.** *A set of sentences Γ is Universal if and only if it is $\Subset$-fragile, i.e. iff $V_M(\Gamma) = 1$, there is an extension of M, N such that $V_N(\Gamma) = 0$.*

Hill writes:

> Prior's proposal seems to depend on the claim that a sentence's classification as ethical or non-ethical is invariant and context-independent. But there is another possibility. . . . Perhaps . . . a sentence's classification is sometimes variable and context dependent. (547)

Hill, S. (2008). 'Is'-'Ought' derivations and ethical taxonomies. *Philosophia*, 36:545–566

finds an earlier variant in . On Humberstone's reading, the classification of mixed sentences is *world relative*—it depends on what other contingent facts obtain. If D is true, then $\neg D \rightarrow N$ will be descriptive; if D is false, the same conditional is normative. Independently of a world to fix the value of D, the conditional does not yet have a status. Relative to a world, the descriptive/normative taxonomy is a dichotomy.

Humberstone, I. L. (2019). Recent thought on *is* and *ought*: connections, confluences, and rediscoveries. *Journal of Applied Logics*, 6(7):1373–1446; and Karmo, T. (1988). Some valid (but no sound) arguments trivially span the 'is'-'ought' gap. *Mind*, 97(386):252–257

[2]

— There are sets of sentences which are Universal, even though they have no Universal members, for example:

$$\{Fa \to \forall x Gx, Fa\}$$

— Any model which makes this set true must make Fa true, and so $\forall x Gx$ true as well. But any such model can be extended to one which makes the set false by adding an element which is "$\neg G$". Hence the set is Universal. However Fa is not Universal, and $Fa \to \forall x Gx$ is Neither (and so not Universal).[3]

— Moreover, we can prove that Universality has a kind of 'infectiousness'— if you add a Universal sentence to a set, it always makes the whole set Universal:

Proof 2 (Infectiousness: If Γ is a set of sentences, and ϕ a Universal sentence, then $\Gamma \cup \{\phi\}$ is Universal.)**.** *Suppose* $\Gamma \cup \{\phi\}$ *is true in model M (recall that unsatisfiable sets are trivially* $\Subset$*-fragile.) Then M makes* ϕ *true and since* ϕ *is Universal, there is an extension of M, N, which makes* ϕ *false. N makes* $\Gamma \cup \{\phi\}$ *false too and so the set is universal.*

[3] Another such set is:$\{Fa \vee \forall x Gx, \neg Fa\}$.

The Universal Barrier Theorem (fragility version)

A striking consequence of defining *Universal* with $\Subset$-fragility is that it makes a Particular/Universal barrier easy to prove.

Particular/Universal Barrier. Let Γ be a Particular, satisfiable set of FOL-sentences and ϕ an FOL-sentence such that $\Gamma \cup \phi$ is Universal. Then $\Gamma \nvDash \phi$. □

Proof. Let M be a model that satisfies Γ. Either M makes ϕ true, or not. If not, then M is a counterexample and $\Gamma \nvDash \phi$. If so, then M makes $\Gamma \cup \{\phi\}$ true and since (by assumption) this set is Universal, there is a model N, $M \Subset N$, which makes $\Gamma \cup \{\phi\}$ false. Now Γ is Particular, so N makes Γ true. So it must be ϕ that is false in N, and again we have a counterexample. Either way: $\Gamma \nvDash \phi$. □

This barrier theorem does not have quite the expected form; rather than talk of the Universality of the conclusion, it talks of the Universality of the union of the premises with the singleton of the conclusion—what we might call the Argument Set: the set of all sentences in either the premises or conclusion. However, recall that we showed that Universal sentences will 'infect' sets to which they are added. From that and theorem , the following weaker barrier theorem also follows:

Particular/Universal barrier (weaker). If Γ is satisfiable and Particular and ϕ is Universal, $\Gamma \nvDash \phi$. □

This weaker theorem has the form we expect, but it also has the defect of remaining silent about Prior-disjunctions and Vranas-conditionals (since these are not Universal on their own)—and that makes it *too* weak. It fails, for example, to say anything above why $Fa \vDash Fa \to \forall x Gx$ cannot be valid. The stronger barrier remedies this defect, and entails the snappy barrier.

The Past/Future Barrier

The name *Hume's Law* is usually reserved for the *is/ought* barrier, but Hume's endorsement of the past/future barrier is also quite 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, 1748, 4.21/37)

What we would need to apply the same strategy in this case:

1. *set up a suitable logic*. To *prove* a barrier, we need a well-defined entailment relation. We get this by specifying a formal language, along with enough model theory to define entailment. At this stage we will accept a lot of help from Prior himself.
2. *identify a binary relation* on the models. This plays the role that extension played in the particular/universal barrier.
3. use the relation to *define the premise and conclusion classes*
4. *formulate the barrier* using the definitions
5. *prove the barrier*

Introduction to Tense Logic

Primitive expressions

sentence letters:	p, q, r, p_1, p_2 etc.
logical connectives:	$\neg, \wedge, \vee, \rightarrow, \leftrightarrow$
tense operators:	$\mathcal{F}, \mathcal{G}, \mathcal{P}, \mathcal{H}$

Sentences

1. If ϕ is a sentence letter, then ϕ is a sentence.
2. If ϕ and ψ are sentences, then
   (a) $\neg\phi$ is a sentence,
   (b) $(\phi \wedge \psi)$, $(\phi \vee \psi)$, $(\phi \rightarrow \psi)$, and $(\phi \leftrightarrow \psi)$ are sentences,
   (c) $\mathcal{F}\phi$, $\mathcal{G}\phi$, $\mathcal{P}\phi$, and $\mathcal{H}\phi$ are sentences.
3. Nothing else is a sentence.

Informally, we can think of $\mathcal{F}\phi$ as saying *at some future time* ϕ, $\mathcal{G}\phi$ as saying that *at all future times* ϕ, $\mathcal{P}\phi$ as saying *at some time in the past* ϕ and $\mathcal{H}\phi$ as saying *at all past times* ϕ. To give the official truth-conditions for these expressions, we need some model theory.

Models

A TL model is an ordered quadruple,

$$(T, <, n, I)$$

where:

1. T is a non-empty set (the set of *times*),
2. $< \subseteq T^2$ (the *earlier than* relation),
3. $n \in T$ (the model's 'now' or present)
4. I is an interpretation function mapping pairs of sentence letters and times into the set of values $\{1, 0\}$, where 1 represents *true*, and 0 represents *false*.

Here's a diagram for a TL model:

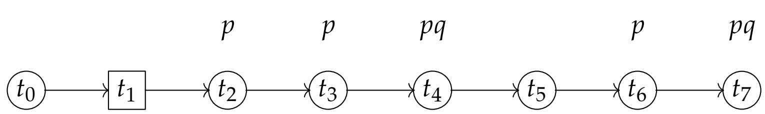

Figure 1: An example of a TL model

This first model is simple, but unrealistic; hopefully there are more than eight times? Perhaps time is infinite, or dense. We can't draw a fully explicit picture for models which represent those things, but I will sometimes indicate additional unpictured points to the right or left with a line of dots, like in Fig. 2:

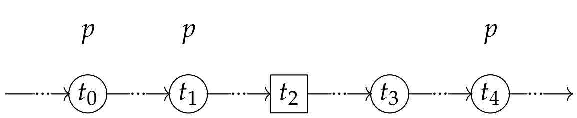

Figure 2: Another TL model

We can model assumptions about the structure of time by imposing different restrictions on $<$. We will assume transitivity, i.e. for all times $u, v, w \in T$,

$$\text{if } u < v \text{ and } v < w, \text{ then } u < w$$

Are times densely ordered by $<$? If so we might adopt: for all times $u, v \in T$,

$$\text{if } u < w \text{ then there is a time } t \text{ such that } u < t \text{ and } t < w$$

Is time allowed to branch, so that one point may lead to multiple futures? If not, we might adopt: for all times $u, v, w \in T$,

$$\text{if } u < v \text{ and } u < w \text{ then } v < w \text{ or } v = w \text{ or } w < v$$

Placing more restrictions on $<$ tends to result in stronger logics, and placing fewer, weaker ones. Here we will use a relatively strong tense logic for the purposes of illustration.

So let $<$ be transitive, anti-symmetric, R-total, L-total, R-extendible, L-extendible, and dense. That is, for any times $u, v, w \in T$:

if $u < v$ and $v < w$ then $u < w$	(transitivity)
not ($u < v$ and $v < u$)	(anti-symmetry)
if $u < v$ and $u < w$ then $v < w$ or $v = w$ or $w < v$	(R-totality)
if $v < u$ and $w < u$ then $v < w$ or $v = w$ or $w < v$	(L-totality)
$u < t$ for some t	(R-extendibility)
$t < u$ for some t	(L-extendibility)
if $u < v$ then $u < t$ and $t < v$ for some t	(density)

Next we extend our I-function to a V-function which assigns 1 or 0 to all the sentences in the language.

Truth at a time in a model

Let V be a function from pairs of sentences of TL and times in T into $\{1, 0\}$ such that

1. If ϕ is a sentence letter, then $V_M(\phi, t) = I_M(\phi, t)$
2. $V_M(\neg\phi, t) = 1$ if and only if $V_M(\phi, t) = 0$
3. (a) $V_M(\phi \wedge \psi, t) = 1$ if and only if $V_M(\phi, t) = 1$ and $V_M(\psi, t) = 1$
   (b) $V_M(\phi \vee \psi, t) = 1$ if and only if $V_M(\phi, t) = 1$ or $V_M(\psi, t) = 1$.
   (c) $V_M(\phi \rightarrow \psi, t) = 1$ if and only if $V_M(\phi, t) = 0$ or $V_M(\psi, t) = 1$
   (d) $V_M(\phi \leftrightarrow \psi, t) = 1$ if and only if $V_M(\phi, t) = V_M(\psi, t)$
4. (a) $V_M(\mathcal{F}\phi, t) = 1$ if and only if there is $u \in T$ such that $t < u$ and $V_M(\phi, u) = 1$
   (b) $V_M(\mathcal{G}\phi, t) = 1$ if and only if for all $u \in T$ such that $t < u$, $V_M(\phi, u) = 1$

(c) $V_M(\mathcal{P}\phi, t) = 1$ if and only if there is $u \in T$ such that $u < t$ and $V_M(\phi, u) = 1$

(d) $V_M(\mathcal{H}\phi, t) = 1$ if and only if for all $u \in T$ such that $u < t$, $V_M(\phi, u) = 1$

In the model in Fig. 2, $\mathcal{F}p$ is true at t_2, but so are $\mathcal{F}\neg p$, $\mathcal{F}(p \wedge \neg q)$, $\mathcal{P}p \wedge F(\neg p \vee q)$, and $\neg\mathcal{G}p$.

Truth in a Model

A sentence ϕ is true in a TL-model M (we write $V_M(\phi) = 1$) just in case $V_M(\phi, n_M) = 1$.

We say a sentence is true in a model *simpliciter* just in case it is true at the model's *now*.[4]

[4] So for example $\neg p$, $\mathcal{F}p$ and $\mathcal{P}\neg q$ are all true in the model in Fig. 2.

We will also speak of *sets* of sentences being true in a model:

A set of sentences Γ is true in a TL-model M just in case $V_M(\gamma) = 1$ for all $\gamma \in \Gamma$. (We write $V_M(\Gamma) = 1$.)

Logical consequence

A sentence ϕ is a TL-logical consequence of a set of sentences Γ (we write $\Gamma \vDash_{TL} \phi$) if and only if for all TL-models M, if $V_M(\Gamma) = 1$ then $V_M(\phi) = 1$.

Prior's other counterexample

— An appropriate logic is a prerequisite for proving a barrier theorem; it is what we prove the theorem about.

— BUT it also refines our tools for formulating counterexamples.

— Prior developed the approach to tense logic that we are using here, but he thought that he had used it to *disprove* the past/future barrier thesis.

Prior himself identifies a precursor for tense logic in J.N. Findlay, and a precursor for his counterexample to the barrier in J. M. E. McTaggart.

— He asks us to "consider e.g., the law CpGPp, 'What is so will-always have been so'." (Prior, 1967, 57) Transposed into our notation this is:

$$\vDash p \rightarrow \mathcal{G}\mathcal{P}p$$

— which is easily transformed into a counterexample to the past/future barrier to entailment:

$$p \vDash \mathcal{G}\mathcal{P}p$$

— In other words, p entails *in the future it will always be the case that at some time in the past p.*

— We can also show:

$$p \vDash \mathcal{FP}p$$

$$\mathcal{P}p \vDash \mathcal{GP}p$$

$$\mathcal{P}p \vDash \mathcal{FP}p$$

These also look like counterexamples.

— If we like, we can set the problem up as another dilemma: Suppose $\mathcal{FP}p$ is *future*. Then

$$\mathcal{P}p \vDash \mathcal{FP}p$$

is a counterexample. Suppose instead that $\mathcal{FP}p$ is *past*. Then

$$\mathcal{FP}p, \neg\mathcal{P}p, \neg p \vDash \mathcal{F}p$$

will serve as the counterexample instead.

— Either way we have a counterexample to the past/future barrier. Still, we saved a barrier from Prior in the previous chapter. Let's see if we can do it again.

Future-switching

— We need a relation on the set of TL-models analogous to that of model *extension*.

— We will use *future-switching*, where one model is a future-switch of another if you can get from the first to the second by *changing what happens in the future*.

— A simple implementation of this idea says that one model is a future-switch of another just in case the models are identical except that for any t, $n < t$, and for any sentence letter ϕ, the values of $I(\phi, t)$ may differ.

— Here's a more precise definition:

Definition (Basic Future-Switching ($\curlyvee$)). *Let M, N be tense logic models. N is a basic future-switch of M ($M \curlyvee N$) just in case*

1. $T_N = T_M$
2. $<_M = <_N$
3. $n_M = n_N$
4. *and for all sentence letters ϕ, and $t \in T_N$, such that $t \leq n$,*

$$I_N(\phi, t) = I_M(\phi, t)$$

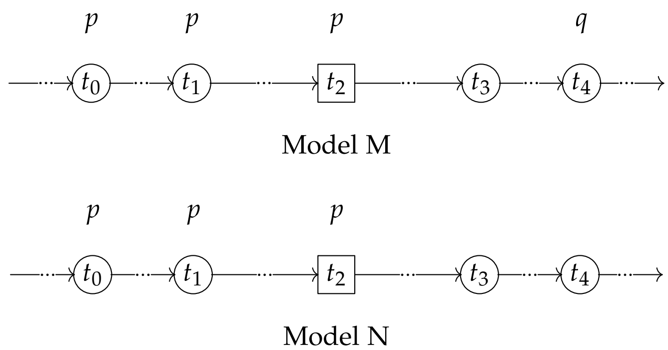

Figure 3: Two models that stand in the basic future-switch relation.

— The truth-values of some sentences will be sensitive to basic future-switching. For example, $I_M(Fq) = 1$ but $I_N(Fq) = 0$, so basic future-switching can make $\mathcal{F}q$ "go false".

— Some other sentences cannot be made false by basic future-switching. If $\mathcal{P}p$ is true in a model, it will be is true in all future-switches of that model.

Taxonomy: Past and Future sentences

— Next we want to use future-switching to classify of the sentences of TL.

— *Past* sentences are future-switch anti-fragile (Υ-anti-fragile) which means that whenever they are true in a model, they are true in all future-switches of that model.

— Similarly we will say that *Future* sentences are future-switch fragile (Υ-fragile) which means that whenever they are true in a model, they are false in at least one future-switch of that model.

— These definitions are promising. They give us four classes of sentence—Past, Future, Both and Neither—and many of the resulting classifications are intuitively right.

— An example is $\mathcal{P}p$. If it is true in a model, there is some $t < n$ where p is true. Since that feature will be preserved through future-switches, $\mathcal{P}p$ will be true in any future-switch of the model. So the sentence is classified as Past, as one would expect.

— Or take $\mathcal{F}p$. If it is true in a model then there is some time t, $n < t$, where p is true. We can construct a future-switch of any such model by changing the value of p at all times later than n to 0. (See Figure 4.) $\mathcal{F}p$ is false in that model, and so the sentence is Future, again, as one would expect.

— This taxonomy also makes it straightforward to prove a past/future barrier theorem:

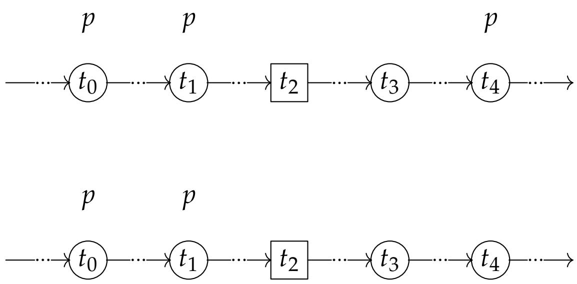

Figure 4: A model in which $\mathcal{F}p$ is true above a future-switch of that model in which it is not.

Past/Future Barrier Theorem. Let Γ be a set of sentences which is Past and satisfiable and ϕ a sentence such that $\Gamma \cup \phi$ is Future. Then $\Gamma \nvDash \phi$. □

Proof. Γ is satisfiable, so there is a model M of Γ. Either M makes ϕ true, or it does not. If it does not, then we have a counterexample to the entailment and so $\Gamma \nvDash \phi$. But if M makes ϕ true, then M makes $\Gamma \cup \{\phi\}$ true and since (by assumption) this set is Future, there is a future-switch of M, N, which makes $\Gamma \cup \{\phi\}$ false. Now Γ is Past (future-switch anti-fragile) and so N makes Γ true. So it must be ϕ that is false in N, and again we have a counterexample. Either way: $\Gamma \nvDash \phi$. □

— A nice feature of this approach is the way it deals with Prior's counterexamples from §:

$$\mathcal{P}p \vDash \mathcal{F}\mathcal{P}p$$

— $\mathcal{F}\mathcal{P}p$ turns out to be Neither (Past nor Future.)

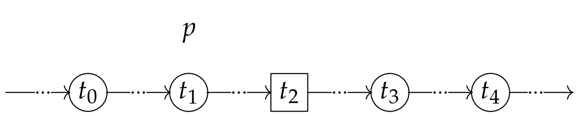

Figure 5: $\mathcal{F}\mathcal{P}p$ is true in this model and no future-switch makes it false.

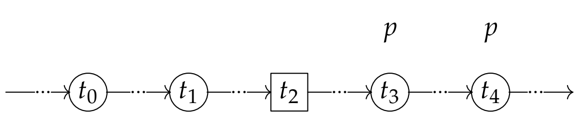

Figure 6: $\mathcal{F}\mathcal{P}p$ is true in this model and some future-switches make it false

A counterexample to theorem needs to be a valid argument, $\Gamma \vDash \phi$, with Γ Past and $\Gamma \cup \{\phi\}$ Future. $\mathcal{P}p \vDash \mathcal{F}\mathcal{P}p$ is not a counterexample because $\{\mathcal{P}p, \mathcal{F}\mathcal{P}p\}$ is not Future. It fails to be $\curlyvee$-fragile because any model of the set is a model of $\mathcal{P}p$, and so has no future-switches that make $\mathcal{P}p$ false, and hence none that make $\mathcal{F}\mathcal{P}p$ false.

Getting more General

— Our proofs of the particular/universal and past/future barriers shared a pattern.
  - In each case we started with a logic—in the sense of a language with a set of models used to define an entailment relation on that language
  - and identified a binary relation (such as model extension, or future-switching) on those models.
  - this relation was used to define sentences, and sets of sentences, which were fragile or anti-fragile with respect to it.
  - Then we formulated the barrier theorem to say that no satisfiable set of R-anti-fragile sentences entailed a sentence if the result of adding it to the premise set would be an R-fragile set.

— Now we can step back to do things in a more general way.

— Instead of starting with a particular logic, (like FOL or TL), we say: suppose you have a logic with a language, L, and a set of models, U.

— And instead of focusing on model-extension, or future-switching, we say: suppose you have some binary relation, R, on U.

— Then we use R to define two sets of sentences of L: those that are R-anti-fragile and those that are R-fragile.

— Finally we use these sets of sentences to formulate and prove a General Barrier Theorem.

Suppose we have a logic, L, where this is to say that we have:

1. a formal language
2. a set of models, U
3. a definition of truth-in-a-model.

1.–3. induces an entailment relation on the language via the following definition:

Definition (Logical consequence in L)**.**
$\Gamma \vDash \phi$ *if, and only if, for all* $M \in U$, *if* $V_M(\Gamma) = 1$, *then* $V_M(\phi) = 1$.

Now let R be a binary relation on U, i.e. $R \subseteq U^2$.

Definition (R-anti-fragility)**.**

1. *A sentence,* ϕ, *is R-anti-fragile on* U, *if and only if for all* $M \in U$, *if* $V_M(\phi) = 1$, *then for all* $N \in U$, *if* MRN, *then* $V_N(\phi) = 1$.

2. *A set of sentences, Γ, is R-anti-fragile on U, if and only if, for all $M \in U$, if $V_M(\Gamma) = 1$ then for all $N \in U$, if MRN, then $V_N(\Gamma) = 1$.*

Definition (*R*-Fragility)**.**

1. *A sentence, ϕ, is R-fragile on U, if and only if for all $M \in U$, if $V_M(\phi) = 1$, then there is at least one $N \in U$, such that MRN and $V_N(\phi) \neq 1$.*
2. *A set of sentences, Γ, is R-fragile on U, if and only if, for all $M \in U$, if $V_M(\Gamma) = 1$ there is at least one $N \in U$ such that MRN, then $V_N(\Gamma) \neq 1$.*

General Barrier Theorem: If Γ is satisfiable and R-anti-fragile but $\Gamma \cup \{\phi\}$ is R-fragile, then $\Gamma \not\models \phi$.
Let Γ be satisfiable and R-anti-fragile and $\Gamma \cup \{\phi\}$ be R-fragile. Since Γ is satisfiable, there is a model, M, that makes it true. If ϕ is not true in M, then M is a counterexample to the entailment, and $\Gamma \not\models \phi$ as required. Suppose then that ϕ is true in M. Then $\Gamma \cup \{\phi\}$ is true in M, but since that set is R-fragile by assumption, there is a model N, such that MRN and $\Gamma \cup \{\phi\}$ is not true in N. But Γ is true in N, because Γ is R-anti-fragile. So ϕ is not true in N, and N is our counterexample to the entailment, and so again, $\Gamma \not\models \phi$. □

References

Hill, S. (2008). 'Is'-'Ought' derivations and ethical taxonomies. *Philosophia*, 36:545–566.

Humberstone, I. L. (2019). Recent thought on *is* and *ought*: connections, confluences, and rediscoveries. *Journal of Applied Logics*, 6(7):1373–1446.

Hume, D. (1975/1748). *Enquiries concerning human understanding and concerning the principles of morals*. Oxford University Press, Oxford. edited by L. A. Selby-Bigge.

Karmo, T. (1988). Some valid (but no sound) arguments trivially span the 'is'-'ought' gap. *Mind*, 97(386):252–257.

Prior, A. N. (1967). *Past, Present and Future*. Clarendon Press, Oxford.

Vranas, P. (2010). Comment on 'Barriers to Implication'. In Pigden, C., editor, *Hume, Is and Ought: New Essays*, pages 260–267. Palgrave MacMillan.