Gillian K. Russell, *Barriers to Entailment: Hume's Law and Other Limits on Logical Consequence*, Oxford: Oxford University Press, 2023, pp. xii+303, ISBN 978-0-19-287473-3 hbk, US$80.00.

1 Introduction

Here are five inferences:

[1] Samara gave money to Donald Trump's campaign. So Samara ought to have given money to Trump's campaign.

[2] Donald is a politician and a narcissist. So all politicians are narcissists.

[3] It has not hit 48° degrees in Melbourne in any year since records have been kept. So it will not hit 48° degrees in Melbourne in 2026.

[4] Yasuo is retiring from his university position this year. So Yasuo has to retire from his university position this year.

[5] It rained on 1/1/2025 in New York. So it rained on1/1/2025 here (*said in Melbourne*).

Each inference is patently invalid. It is often averred that this is so because, in each case, the premise and the conclusion are of different *kinds*, and one cannot validly infer a statement of the conclusion kind from statements of the premise kind.

Picturesquely, one might say that there is a deductive inferential barrier between each of the two kinds. The invalidity of the inferences [1]–[5] are, respectively, standard examples of the supposed barriers of inferences from:

- *is* to *ought* [I/O] ("Hume's Law")
- *particular* to *universal* [P/U]
- *past* to *future* [P/F][1]
- *is* to *must* [I/M]

[1]Or better, as it turns out, *past and present* to *future*.

- *non-indexical* to *indexical* [N/I]

Call the claims that there are such barriers *barrier claims*.

Are such barrier claims correct? Somewhat notoriously, there are well-known putative counterexamples to such claims, which have generated a substantial literature. (We will look at some of these examples in due course.) The main aim of Russell's book is to defend the existence of the barriers in question when they are appropriately articulated.

Undoubtedly there are deductive inferential barriers. Let $\mathcal{L}$ be a formal language, and let $\vdash$ be a consequence relation between formulas of $\mathcal{L}$. Let us assume (as does Russell) that we are dealing with a single-consequence relation. Let $\Gamma_1 \subseteq \mathcal{L}$; let $\Gamma_2 \subseteq \{A : \Gamma_1 \nvdash A\}$. Then nothing of kind Γ_2 can be inferred from premises of kind Γ_1. The question is, then, whether the kinds in question in the barrier claims can be made to line up with such a Γ_1 and Γ_2. Russell claims that they can; and moreover, that they can be seen to do so in a uniform way: there is a construction which shows that the five barriers are special cases of a single general kind.

The book proceeds as follows. After an introduction to the project, Chapter 1 reviews the standard objections to the barrier claims. Chapter 2 discusses P/U; Chapter 3 discusses P/F. Russell's favoured articulation of the barrier will appear only after all five cases have been considered. Chapter 4 is a sort of mid-way staging post. In this she formulates a first stab at the appropriate articulation, the so called General Barrier Theorem. This will turn out not to provide what is required. The next three chapters then discuss the other cases, which motivate a more adequate formulation. Chapter 5 discusses I/M; Chapter 6 brings P/F and I/M together, establishing a problem with the strategy thus far followed; Chapter 7 discusses N/I. Finally, Chapter 8 discusses I/O, before Chapter 9 formulates the final articulation of the uniform barrier, in the shape of the so called Limited General Barrier Theorem.

The final two chapters then take a new turn. So far we have been concerned with formal validity, but many have suggested that some counterexamples to the barrier claims cannot be handled with formal techniques. The last two chapters address the issues here. Chapter 10 discusses how the formal ideas so far used can be modified to apply appropriately to "natural language". Chapter 11 then uses these to establish informal results analogous to those of the formal part of the book.

There is much of interest in Russell's book. Her discussions are careful, thoughtful, and often insightful. There is nothing original about the formal tools she employs, but they are employed appropriately and deftly.

It is not possible to discuss all of the many issues raised by the book: the subject matters of the inferences involved are of too diverse a kind. In the first part of this essay, and until further notice, I will discuss the formal chapters of the book. In the second part, I will turn to the informal chapters.

2 The Limited General Barrier Theorem

I will not discuss all the twists and turns of the book's journey to its final position (which appear to mirror Russell's own journey of several years to reach a position she finds adequate). I will merely consider this position.

The methodology for each of our five boundaries is the same. We start with a formal language in which inferences that cross the relevant boundaries can be expressed. Each of these comes with a standard model-theoretic definition of validity. (In the book Russell assumes, but does not defend, the correctness of a model theoretic account of validity.) All of these are based on classical logic—a somewhat strange position for a self-confessed logical pluralist.[2] These are: deontic logic [I/O]; first-order logic [P/U]; any normal modal logic [I/M]; a standard tense logic (with transitive temporal relation) [P/F]; a logic for indexicals [N/I]. The logics can be combined when appropriate. Indeed Chapter 9 puts them all together. All of these logics are familiar logical fare,[3] with the exception of the last, which is a simplified version of Kaplan's *Logic for Demonstratives*.[4]

Russell calls interpretations for the languages *models*. This is standard enough, though I think it better to use that term for interpretations that make a given set of sentences true. However, I will follow her in her nomenclature. Given a model, M, for the appropriate logic, $V_M(\phi) = 1[0]$ means that the value of ϕ in M is 1[0]. (In the logics with world semantics, a world-independent truth value is obtained by having a designated "actual"

[2]See Russell (2008). In fact, given standard motivations for relevant logic, it might be thought that this would be particularly appropriate for the project—though Russell rejects this (p. 38 ff). For an exploration of what can be done using relevant logic, and a brief discussion of Russell's rejection, see Weiss (202+).

[3]For example, they can all be found in Priest (2008).

[4]Kaplan (1978).

world.) If Γ is a set of sentences, then $V_M(\Gamma) = 1$ means that $V_M(\phi) = 1$ for all $\phi \in \Gamma$, and $V_M(\Gamma) = 0$ means that $V_M(\Gamma) \neq 1$, that is $V_M(\phi) = 0$ for some $\phi \in \Gamma$.[5]

In each of our cases, the class of appropriate models is furnished with a binary relation, R. We then have the crucial definition (p. 203 f):

- A sentence ϕ is *R-breakable* if there is at least one pair of models M, N, such that MRN and $V_M(\phi) = 1$, but $V_N(\phi) = 0$.

- A set of sentences Γ is *R-breakable* if there is at least one pair of models M, N, such that MRN and $V_M(\Gamma) = 1$, but $V_N(\Gamma) = 0$.

So, note, Γ is *not* R-breakable (*R-unbreakable*) iff for every M and N such that MRN and $V_M(\Gamma) = 1$, $V_N(\Gamma) = 1$.

In each case, the premise-kind of statements is identified as the class of statements that are R-unbreakable; the conclusion-kind of statements is identified with the class of statements that are R-breakable.

The statement of the Limited General Barrier Theorem is then as follows. (I rephrase for the purpose of clarity.[6]) Given the relevant set of models, and the relation R on them, then $\Gamma \not\models \phi$ if:

(A) Γ is R-unbreakable

(B) ϕ is R-breakable

and:

(C) There are models, M and N, such that MRN, $V_M(\Gamma) = 1$, and $V_N(\phi) = 0$.

Condition (C) rules out putative counterexamples to a barrier claim where Γ is not satisfiable, or where ϕ is always satisfied. It seems very natural to regard these as special cases, and to take them out of the picture. But (C) does more than that, as we will see.

---

[5]The last clause is not made explicit in the book, but was clarified for me by Russell in correspondence.

[6]What Russell actually says (p. 206) is: *if the sentence ϕ is R-breakable, but the set of sentences, Γ, is not then $\Gamma \not\models \phi$ unless each model M which makes Γ true is such that all models N to which is it R-related are also models of ϕ.* Assuming that 'unless' is truth functionally equivalent to 'or', this is ambiguous between sentences of the form $((A \wedge B) \rightarrow C) \vee D$ and $((A \wedge B) \rightarrow (C \vee D))$. Fortunately, these are both equivalent (in classical logic) to $(A \wedge B \wedge \neg D) \rightarrow C$.

Anyway, the proof of the Theorem is almost trivial. Given Γ and ϕ, choose models M and N as given by (C). Since Γ is R-unbreakable (by (A)), $V_N(\Gamma) = 1$. So N is a counter-model to the inference.[7]

In other words, given conditions (A), (B), and (C), there can be no counterexamples to the barrier claim. If there is anything interesting about a putative counterexample, then, this can only be why, exactly, it fails. Of more interest is the definition of various Rs.

3 The Five Barriers

So let us look at our five cases in the light of this. In each case, we will look at the definition of R, and why, given it, one of the standard counterexamples to the barrier claim fails.

3.1 *Particular/Universal*

First, the most straightforward case: *P/U*. In this case, MRN iff N is the same as M except that its domain extends that of M. So constants do not change their denotations, and objects in the domain of M remain in the extensions of whatever predicates they were in M. In a clear sense, no particular facts have changed (though there may be some new ones); this makes it natural to think of R-unbreability as tracking the preservation of particularity. And since we have some new objects to play with, some general facts may well change. This makes it natural to think of R-breakability as tracking generality.

As a putative counterexample, consider $Fa \vee \forall x Gx, \neg Fa \vdash \forall x Gx$. The conclusion is R-breakable; but so is the set of premises. For let M be a model which makes both premises true. Then Fa is false, and $\forall x Gx$ is true. Let N be a model which is the same except that some new object does not satisfy G. Then MRN, and the first premises is false. So this is not a counterexample to the barrier theorem.

[7] It is worth noting that the proof is entirely independent, not not only of the definition of R, but even of the logic for which the model theory is deployed.

3.2 *Is/Must*

Next, let us look at the I/M case. Here we have a possible-world semantics to deal with. In this case, MRN iff all the worlds in M are in N, the "base worlds" are the same, and for any world, w, in M, the value of any atomic sentence is the same in N. Generally, in a world semantics, atomic sentences are thought of as factual. So R-unbreakability can naturally be seen as tracking what is the case. The fact that N adds new possible worlds, and so new possibilities, makes it natural to see R-breakability as tracking what must be the case.

As an example, let us consider the inference $p \models \Box\Diamond p$. If the accessibility relation of M is not symmetric, the inference is not valid. If it is, let M be a model such that $V_M(p) = 1$; so in any N such that MRN, $V_N(p) = 1$, and $V_N(\Box\Diamond p) = 1$. Hence condition (C) is not satisfied.

3.3 *Past/Future*

Turning to the P/F case, we are dealing with a tense logic. MRN if N is the same as M, except that, where the temporal ordering is $<$ and @ is the distinguished world (the present), the set of times $\{t : t > @\}$ may vary, and for any $t \leq @$ things are the same. The fact that N changes future and only future facts makes it natural to take R-breakability to be future, and R-unbreakability to be past (or better past or present).

Turning to a possible counterexample, consider the inference $p \vdash \mathbf{GP}p$.[8] Let M be a model where p holds at @. By the truth conditions of **G** and **P**, so is $\mathbf{GP}p$. So condition (C) fails.

3.4 *Is/Ought*

The I/O case deals with a deontic logic. MRN holds if N is the same as M except that the deontic accessibility relation is different. This does not change the facts at any world, so R-unbreakability plausibly tracks what is the case. However, changing the accessibility relation changes the deontically possible worlds, and so R-breakability tracks what ought to be the case.

[8] Russell uses calligraphic script for both tense operators and deontic operators. I boldface the tense operators for clarity.

In the light of this, let us consider the following putative counterexample: $p \vee Oq, \neg p \vdash Oq$. The conclusion is obviously R-breakable. But so is the set of premises. For let M be a model where the premises are true at @. Then p is false there, and at every world accessible to @, q is true; but by changing the accessibility relation, and taking an N such that there is at least one world where q is false, we can make Oq false at @. Hence, the set of premises is breakable, and we do not have a counterexample.

3.5 *Non-Indexical/Indexical*

The final sort of case concerns N/I. MRN holds if N is the same as M, except that the agent and the place—the relevant indexicals in the language—may change. So R-breakability shows that truth depends on the denotations of the indexicals in the sentence, and R-unbreakability shows that it does not.

Let us now consider a putative counterexample: $\forall x Px \vdash P\mathfrak{i}$, where $\mathfrak{i}$ is an indexical referring to the agent. The premise is not R-breakable, but the conclusion is. Yet the inference is valid. However, in this case clause (C) fails. If M is any model in which the premise is true, the conclusion is true whatever the denotation of $\mathfrak{i}$ is. So there is no N such that MRN and the conclusion is false in N.

It certainly appears, then, that Russell's strategy provides a reasonably natural way of characterising the relevant barriers, and disarms the counterexamples—as it must. However, that is not an end of the matter.

4 Categorising the Kinds

In the inferences of §1, that the premise and conclusion belong to the relevant kinds of classes is pretty obvious. However, these classes are clearly vague, and in many cases, intuition provides no help in making matters precise. For example, if A is 'Donald is a narcissist', B is 'Yasuo is retiring', are $A \vee \Box B$ or $A \rightarrow \Box B$ *is*s or *must*s? Intuition gives no clear answer. Being vague, the distinctions can be precissified in a number of different ways. Indeed, Russell discusses many such possible precissifications in the journey to her final position.

For a start, Russell's modal case is a construction that allows for different things to be possible; her deontic case allows for different things to be permissible. In the modal case this is achieved by adding more worlds; in the deontic case, this is achieved by changing the accessibility relation. These are both perfectly natural, but each strategy could be reversed (or the two cases could be treated in the same fashion). Though these are different strategies, they both fit into Russell's general framework which delivers the Barrier Theorem.[9]

Different natural precisifications could be more radical, however. For example, in the P/U case, we could take a particular sentence to be one which is quantifier-free, and a universal one to be one with quantifiers. This might be thought to be odd, since it makes $\exists xPx$ universal. But this is equivalent to $\neg\forall x\neg Px$, and so still tells you something about a universal situation. After all 'Mary is not happy' tells you something about Mary just as much as 'Mary is happy'. Given this approach, there is no natural barrier between particular premises and universal conclusions, since $Pa \models \exists xPx$.

As with most precissifications of vague notions there seems to be no uniquely correct right way to precissify here. Sometimes Russell seems to suggest otherwise (p. 135):[10]

> In general, the danger is often less that the barrier will fail, but that the taxonomy will turn out not to fit with *the classificatory barrier* we seek.

But I take it that her considered view is that this is not the case: in principle one could go many ways. Thus, towards the end of the book, in discussing the informal I/O barrier, she says (p. 283):

> Someone might object that I have simply assumed that my taxonomy is correct... but what really I have assumed is that I am entitled to specify the thesis I wish to defend. If someone wishes to maintain that there are some interpretations of *normative* on which Hume's law fails, I have no quarrel with them. My thesis is that there is a plausible interpretation of Hume's Law on which it is true and on a par with other—uncontroversially true—barrier theses in philosophy.

[9] In the final section of the formal part of the book (§9.7), in what appears to be something of an afterthought, Russell does note that in the *is/ought* case one could deploy essentially the same definition of *R* as in the modal case, though, she says, she finds this 'less natural'.

[10] My italics. All italics in quotations from Russell are hers unless otherwise noted.

Okay. Then Hume's law can be formulated in different ways. In some forms it may hold; in others it may not. But is the plausible interpretation of the I/O barrier in question on a par with others (let alone uncontroversially true others)? Certainly not obviously, since there are very clear differences between the cases.

Prima facie, it is impossible to get an *ought* from an *is*. But it is just as impossible to get an *is* from an *ought*. Notoriously, the fact that someone ought to do something does not imply that they do do it. Taking our cue from the P/U case, one may accommodate both deontological barriers in the same way, by taking an *ought* sentence to be any sentence that contains a deontic operator, and an *is* sentence to be one that does not. (And both barriers are then explained simply by the fact that both $Op \vdash p$ and $p \vdash Op$ are deontically invalid.) On the other hand, it may not be possible to get a universal conclusion from a particular premise, but it is obviously possible to get a particular conclusion from a universal premise. So this barrier is not symmetric. Symmetry *vs* non-symmetry wouldappear to betoken two different kinds of barriers.

5 Characterising the Barriers Between Kinds

But it is at least the case that Russell has shown that there is one way of precissifying the various notions so that barriers are structurally the same. Or has she? Trouble starts to appear when one looks at the precise formulation of the Barrier Theorem and its proof.

For a start, recall that the precise nature of R plays no role in the proof of the Barrier Theorem; and because of this, its formulation appears to over-generate. Let $|M|$ be the cardinality of the domain of a first-order model, and take MRN to be $|M|>|N|$. Then breakability appears to track the notion of being about fewer objects; and unbreakability tracks the notion of being about more (or more precisely, at least at least as many) objects. The Barrier Theorem holds for this notion of R, but no one has ever suggested that there is an inferential barrier from *more* to *less*—and for good reason. Obviously, if you can establish something about all of a bunch of objects, it follows of each of some smaller subset of them.

Further problems appear when one looks at the problematic role of condition (C) in the proof of the theorem. Up to the end of Chapter 6, Russell's preferred way of characterising the distinction between the kinds of statement on each side of a putative barrier is not in terms of breakability, but in terms of fragility (a related but different notion[11]). But on putting modal and tense logics together, this allows counterexamples to the P/F barrier of the form $\Box p \models \mathcal{F}p$. We then have (p. 149):

> Entailments like this were our main reason for not formulating the barriers in terms of breakability rather than fragility in the first place. Fragility made it easy to prove barriers. But with breakability there were counterexamples, like $\neg p \models p \rightarrow \mathcal{F}p$.[12] Our best hope for a barrier formulated in terms of breakability, then, is a *limited* barrier that says something like:
>
> If ϕ is R-breakable and Γ is not, then $\Gamma \not\models \phi$ *unless* C
>
> where C is to be replaced with some condition on Γ and ϕ.

In other words, clause (C) is there explicitly to rule out counterexamples. As a way of establishing the existence of barriers, then, it is explicitly *ad hoc*.

Russell comes close to admitting this. When discussing informal counterexamples to the I/O barrier, she says (p. 274 f; my italics):

> Counterexamples which fail on the last point [GP: condition (C)] have a more interesting status than the others. After all, *they really are valid arguments with descriptive premises and normative conclusions*, and so genuine counterexamples to the initial hyperbolic statement of the Normative Barrier Thesis. Perhaps, then, as defenders of the Limited Normative Barrier, our attitude to these arguments should be less, "this is where you went wrong" and more, "you were right; see how much we've learned".

The italicized phrase comes perilously close to admitting that invoking clause (C) is really an admission that the barrier claim is an artifice.

In fact, matters are worse than this. A glance at the proof of the Barrier Theorem shows that it is condition (C) that is doing all the real work. Indeed, the condition comes precariously close to itself saying that the inference $\Gamma \vdash \phi$ is invalid.

[11] ϕ is fragile if there are models, M and N, such that MRN, $V_M(\phi) = 1$, but $V_N(\phi) = 0$.

[12] GP: The text says $p \rightarrow Fp$, but I take it that this is a typo.

Note that clause (B) is not used in the proof. So the existence of the barrier is not even dependent on the nature of the conclusion. If the failure of the inference holds for conclusions of *any* kind, it has nothing to do with the kind relevant to the supposed barrier at issue. It's no more a barrier between the *is*s and the *ought*s than it is one between the *is*s and the *must*s—or indeed, the *is*s and the *is*s. And the same goes for all the other putative barriers.

Russell addresses the issue with respect to one putative barrier, the I/O barrier. If you are going somewhere it is good advice to say 'If you take the train, make sure you catch the right train', but that's so obvious as to be irrelevant to most conversations. However, if you are travelling from Flinders St to Burnley, there is a special danger, since some of the trains from there go to Burnley and some do not. So the advice 'If you take the train from Flinders St to Burnley, make sure that you catch the right train' though weaker, is relevant. Similarly (p. 192):

> with descriptive conclusions, there is no danger of s-shifting [GP: that which could give rise to the change of truth value relevant to *R*-breakability] leading to a sentence having the wrong truth-value. With normative sentences there is. So we say: no normative sentence follows from descriptive sentences unless all the s-shifts of the models of the premises are models of the normative conclusions—even though the stronger claim is true, because the claim that makes specific reference to normative conclusions highlights the fact that these conclusions are s-shift breakable.

So we have a speech-act reply. To say that there is no violation of the barrier—whatever the conclusion—would be a violation of Grice's maxim of relevance.[13] Now, what is relevant depends on the context, and you might well think that to say that that counter-examples to the barrier are ruled out for *all* sentences, not just normative ones, is highly relevant to the present context. But in any case, a speech-act reply is really beside the point. We are talking about what kinds of inference are valid or invalid. Conversational implicature is a red herring.

[13]See, e.g., Grandy and Warner (2021), §3.

6 Informal Models

Let us now turn to the last part of the book. Though it comprises only two chapters, I think that this is, in many ways, the most interesting and original part of the book. It certainly deals with issues that go well beyond the barrier issue, though it is hardly possible to discuss all of them here.

Many counter-examples to barrier claims appear to be of a kind different from those we have been considering. Consider, for example:

- The fridge is empty. So there is nothing in the fridge.
- Jabari promised to return the money she borrowed. So Jabari ought to return it.

Prima facie, these appear to violate the P/U and I/O barriers, respectively, though they are not formally valid since they turn on the meanings of the words 'empty' and 'promise'. Similar apparent counterexamples beset the other three putative barriers.

The point of last two chapters of the book is to defend the barrier claims against such examples. The method for doing this is to modify the formal machinery of the earlier parts of the book to apply to natural language, and apply it in an analogous way. As Russell puts it, what is required is the notion of (p. 218):

> an *informal model*, something that is to natural language as formal models are to formal languages.

Now, one might naturally be skeptical that there is such a thing, since the application of model-theoretic techniques work because formal languages have crucial features that natural languages lack. They have a precise grammar; natural languages have no clear grammar—or perhaps better, people make sense of natural language sentences that violate nearly any putative grammatical rule.[14] Ambiguity is rife in natural language, as are metaphor and idiom. These things are absent from formal languages. At the very least, then, natural language has to be sanitised/regimented before model theoretic techniques can be deployed. Russell does not discuss these matters, though she could claim some notable philosophers, such as Richard Montague,[15] on her side—though she does not mention him.

[14]Read Joyce's *Fingers Wake*, or poems by e. e. cummins.

[15]Montague (1970).

Assuming that these matters can dealt with, a glance at the definitions of *breakability* and the Limited General Barrier Theorem show that informal models must:

- Make it possible to assign a truth value (at least relative to a context) to each sentence.
- Allow for a variation of truth values
- Deliver the appropriate *R* relations.

Let us leave the third of these aside for the moment, and consider the first two. These are the concerns of Chapter 10.

In the formal models that Russell has been using in the first part of the book a model posits (a) a domain of objects, (b) a denotation for each name, and (c) a denotation (extension) for each predicate. (I consider only first-order logics, to keep things simple.) If there are "worlds" in the model, it also specifies the set of these, an accessibility relation if one is required, and (a), (b), and (c) for each of the worlds. These features are used to give a recursive definition of truth (at the base world) for each model. An informal model should do the same.

Russell endorses the thought that there is a recursive truth definition for a natural language such as English (when suitably sanitised, I presume), though producing such a thing is a highly non-trivial exercise (p. 240 f):

> [this] is an ambitious project in natural language semantics, requiring expert training that goes beyond most speakers.

Actually, the compositionality of natural language, on which a recursive truth definition depends, is a highly contentious matter.[16] But as Russell goes on to point out, in effect, it is not necessary to assume that there is such a thing. All that we require is that speakers have a good general ability to evaluate the truth of various sentences in various circumstances—presumably by using their imaginations. There may be controversial cases, but doubts that arise here are equally likely to arise concerning the correctness of the clauses of a recursive semantics.

In any case, with or without recursion, the features provided by the model can vary from model to model, producing different truth values for

[16] See Gendler-Szabó (2020).

one and the same sentence. How are we to understand such variation? Somewhat notoriously, the variation in formal models tends to be thought of in two quite different ways.[17] It can be thought of as giving the non-logical-constants different meanings, so as to explain the validity of $\forall x Fx \vdash Fa$, whatever F and a mean. Alternatively it can be thought of as variations in the reality that the language describes, as one thinks of various models of the axioms of group theory. Call the first sort of variation *semantic* and the second *Worldly*. (I capitalise the word to make it clear that this has nothing to do with the worlds of models.)

Neither way of thinking suits Russell's purpose. The variation in the P/U barrier just involves adding more objects to the domain. This has nothing to do with semantics. On the other hand, the fact that $a = b$ is true in some models and false in others seems to make little sense if we are just thinking about a, that is, b, having different properties. The change required here seems semantic rather than Wordly.

In the case of informal models, Russell suggests, both kinds of variation are to be allowed. She calls this the *combination* view. And there seems something very right about this. The truth (or otherwise) of a natural language sentence is determined by both semantic and Worldly factors. 'Canberra is the capital city of Australia' is true, in part because of the meanings of 'Canberra' 'capital city' etc. But it is also true in part due to certain historical and geographical facts about where it was decided to put the federal capital when Australia federated in 1901.

However, this does leave Russell with something of a problem. The putative informal counterexamples to the barrier claims above were due to the meaning of certain (English) words. Of course the inference can be made invalid if we are allowed to change the meanings of the crucial words (e.g., taking *empty* to mean *full*). A proponent of the counterexample is obviously not going to be impressed by this.

Russell's response is to distinguish between two dimensions of meaning. The first she terms *conditional meaning*. This is the meaning a word has in virtue of the 'conditions' of use, such as linguistic rules employed by speakers (p. 237). The other she terms *environmental meaning*. This is the 'part of the world' picked out by the word (p. 237). These need not be the same, as indexicals show. 'I' is used by a speaker to pick out themself. That is its conditional meaning. But when used, its environmental meaning is

[17] As discussed most prominently by Etchemendy (1999).

the person themself. It is natural to think of this distinction as that between *Sinn* and *Bedeutung*, in Frege's terms—though Russell does not suppose (as does Frege) that a linguistic item must have both dimensions. The variation of meanings allowed are variations of environmental meaning, not conditional meaning. The cheap way of invalidating the counterexample inferences varies the conditional meanings of the words. (So you cannot change the sense of *empty* to mean *full*, though you can change what is in the fridge.)

One might worry about the rationale for thus restricting meaning-change. But it is clear that some restriction on meaning change is necessary. If one can change the meaning of words arbitrarily, then virtually any inference could be made invalid. Thus, 'some Germans are women' entails 'some women are Germans'. But this would no longer be the case if we were free to take *some* to mean *all*. A traditional move is to allow only non-logical-constants to change their meanings; but I think trying to distinguish between logical constants and non-logical-constants in any principled fashion is now regarded (correctly) as something of a lost cause. Russell's strategy is more general and seems more robust.

Of course, there are problems to be faced. Frege is often criticised on the ground that the sense of a word varies from person to person.[18] In §10.6 Russell, drawing on examples from Waismann and others, worries about an analogous problem for conditional meaning. Whilst admitting that there can be such variability, she finds, for various reasons, that the worry is overplayed. In particular, she suggests that the conditional meanings of the words involved in the putative informal counterexamples to the barrier claims do not have this variability. For example, it is simply the case that 'to satisfy *empty*, an object must have nothing inside' (p. 247). Perhaps; but what counts as nothing is highly variable. If I say that the fridge is empty, I do not mean that there is no air inside. When someone talks about barometers and says that the vacuum tube is empty, they mean exactly this. When I say that the lecture theatre is empty, I do not mean that there is no furniture inside. When the removalists say that the room is now empty, they do.

Moreover, and worse, this sort of variability can affect the validity of putative barrier claims. Take the term 'fruit'. If you ask a botanist, they will tell you that tomatoes are fruit. But in a greengrocer's shop, you will

[18]For discussion, see May (2006).

find tomatoes in the vegetable section, not the fruit section. There is no sense in which one is right and the other wrong. These institutions just use the word differently. So is the inference *Linqi was asked to buy some fruit and bought only tomatoes. So Linqi did something she ought not to have done* valid or not? It starts to look, at the very least, as though certain barriers may be relative to somewhat arbitrary features of linguistic sub-cultures.

7 Informal Barriers

Let us set these matters aside, and turn to our third bullet point of the last section. This takes us to Chapter 11.

The statement of the informal Limited General Barrier Theorem and its proof are essentially as before, and so have the same issues as the formal case. Moreover, as before, the proof does not depend on the definition of the various *R*s; but the *R*s are important for other reasons.

Given an informal model, the natural informal correlates of the formal *R*s suggested by Russell are:

- P/U: adding new objects, with their properties, to the domain
- I/M: adding new possibilities to those under consideration
- P/F: changing what happens at future scenarios, or changing the possible structure of future times
- I/O: changing what situations are normatively possible
- N/I: changing the denotations of indexical expression

We may ask two questions concerning these.

First, given the "theorem", there can be no counterexamples, but we may ask how the *R*s diagnose the failure of the mooted counter-examples. This is not the place to consider all of these, but let us look quickly at the two examples I noted in the last section:

- The fridge is empty. So there is nothing in the fridge.

Let *M* be an informal model where both premise and conclusion are true. Let *N* be an informal model which is exactly the same, except that there is some cheese in the fridge. Then both are false in this. So both premise and conclusion are *R*-breakable; that is, universal. This is not a counterexample.

- Jabari promised to return the money she borrowed. So Jabari ought to return it.

Let *M* be an informal model where both premise and conclusion are true. Let *N* be an informal model which is the same, except that there is no moral obligation to keep promises—perhaps one where no one ever takes an utterance of 'I promise to...' seriously. Then the conclusion is false. Now, you might say that we knew all along that the inference was not valid. Maybe Jabari did promise, but then she died, so one can no longer attribute an obligation to her. But if the inference is valid, and the conclusion is false, so is the premise. Whatever it takes for Jabari to have made a promise, these conditions are not satisfied. So the premise is *R*-breakable, and there is no counter-example to the Barrier Theorem.

Second, and perhaps of more importance, is the question of whether, in fact, we have an adequate grasp of the way that the various *R*s are supposed to function. Specifically, do they provide an adequate understanding of truth in the relevant counterfactual situations?

In the case of N/I, this is unproblematic: the change of denotation of an indexical term is a familiar and mundane matter. The P/U case is not much more problematic: we have little trouble envisaging what would happen if there are more things than are taken to be the case. The P/F case is again relatively routine, since whenever we have to decide what to do we consider different possible futures. Maybe it's not so hard to consider the possibility of different structures of future times as well—for example, what it would be like if time came to an end at midnight, December 30, 2025.

The I/M case is more complex, simply because there are many kinds of *must*s. Think of some of the most standard ones. There is logical necessity. It's plausible that what's true at all possible worlds may change by varying what logically possible worlds are available—assuming that there is one correct logic (and we know what it is); but what to make of the matter if one is a logical pluralist (as Russell is) is much less clear. Physical possibility is also not too hard, at least given the assumption that what we presently take to be the "laws of nature" is correct. Metaphysical possibility is harder. Even if we suppose that there is such a thing,[19] what things are metaphysically possible is much more problematic, since there is virtually no consensus about what the "laws of metaphysics" are. Then

[19]Which one may not. See Priest (2021).

there is epistemic necessity. The bounds of epistemic necessity are most unclear; but in any case, normal modal logics of the kind that Russell is using do not play nice with epistemic necessity, for familiar reasons (such as the problem of "logical omniscience").

This does not exhaust the list. There is the *can* of ability, the *can* of opportunity, and so on. The plurality of *must*s and *can*s does not really feature on Russell's agenda. Then of course, there is the *can* of permission. This takes us to the I/O case. For a start, there are as many kinds of oughts as there are kinds of norms. The plurality here *is* on Russell's agenda, and provides the longest and most interesting discussion of the various *R*s. In some cases, getting one's head around norm-change is routine—for example, with the rules of games. It is easy to conceive what moves in chess would be (il)legitimate if one dropped the rule of castling. Similarly with the rules of etiquette. Anyone who has travelled knows well what it is like for these to vary. Much the same goes for legal *musts*.

Perhaps the case that puts up the stiffest resistance to imagining variation concerns moral norms. Can we really get our head around what it would be like for it to be morally permissible to torture human babies for fun? Of course, it is easy to get one's head around what it would be like for someone to have bizarre moral *beliefs*, or *change* their beliefs about what is morally permissible, as Russell discusses insightfully. But unless one is a subjectivist about moral norms (which neither I nor Russell are), that is a different matter.

In summary, we see again that there is much more variety in the different kinds of inferential boundaries, and the sub-kinds within each kind, than appears in Russell's one-size-fits-all picture.

8 Conclusion

As will be clear, I have disagreed with Russell about a number of matters. It remains the case that Russell's book provides the most sophisticated and systematic extant account of the vexed question of barrier claims. Any future discussions of the matter will have to start from her book.[20]

Graham Priest, Departments of Philosophy, CUNY Graduate Center, University

[20] Many thanks for comments on an earlier draft of this review go to Antonella Mallozzi, Yale Weiss, and especially Gillian Russell herself.

of Melbourne, and Ruhr University of Bochum.

References

Etchemendy, J. (1999), *Concepts of Logical Consequence*, Stanford, CA: CSLI.

Gendler-Szabó, Z. (2020), 'Compositionality', in E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/entries/compositionality/.

Grandy, R., and Warner, R. (2021), 'Paul Grice', in E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/entries/grice/.

Kaplan, D. (1978), 'On the Logic of Demonstratives', *Journal of Philosophical Logic* 8: 81–98.

May, R. (2006), 'The Invariance of Sense', *Journal of Philosophy* 103: 111–144.

Montague, R. (1970), 'English as a Formal Language', pp. 189–224 of B. Visentini *et al*. (eds.), *Linguaggi nella Società e nella Tecnica*, Milan: Edizioni di Communita; reprinted as pp. 188–221 of Thomason (1974).

Priest, G. (2008), *Introduction to Non-Classical Logic: From If to Is*, Cambridge, Cambridge University Press.

Priest, G. (2021), 'Metaphysical Necessity: a Skeptical Perspective', *Synthese* 198: 1873–1885.

Russell, G. (2008), 'One True Logic?', *Journal of Philosophical Logic* 37: 593–611.

Thomason, R. H. (ed.) (1974), *Formal Philosophy*: *Selected Papers of Richard Montague*, New Haven, CT: Yale University Press.

Weiss, Y. (202+), 'A Relevant Framework for Barriers to Entailment', *IfCoLog Journal of Logics and Their Applications*, forthcoming; currently accessible at: https://philpapers.org/archive/WEIARF-2.pdf.