

Generics: Inference & Accommodation

Greg Restall* · Philosophy Department, The University of St Andrews

JANUARY 2023

Generic claims—such as *Birds fly*, *Men are violent*, and *Mosquitos carry Ross River Fever*—seem pervasive across human thought and talk. We use generic claims to express our understanding of the world around us and our place in it. These generic claims are useful even though they admit exceptions. We can agree that birds fly, even though emus don't. Mosquitos carry Ross River Fever (hereafter, RRF), but not those in Africa. And you can agree that men are violent while conceding that not *all* men are, or even that most are. Generics remain important in our thought and talk in the presence of these counter-instances. Generic claims express rules of thumb, ways to see the world around us, and they provide heuristics for navigating that world. Generics also play a significant role in our maintaining the boundaries of social kinds, and in our attempts to shift those boundaries.

Exactly what does it take for a generic claim of the form *Fs are G* to be *true*? This is a matter of immense theoretical interest, on which there is very little agreement [23, 26, 29, 30, 32, 46]. Beyond theory, things little better. Arguments about contested generic claims can produce much more heat than light. When the topic of men's violence against women is raised, it is a common refrain to hear the defensive retort “not *all* men,” as if that were an objection to the claim of male violence. It seems clear that generics play a significant role, particularly, in the ideologies of our social worlds, of characterising different social kinds and expressing our default orientations toward them [18, 25], and towards ourselves as members of those kinds.

In this paper, I aim to explain the connection between generic claims and our practices of *inference* and *explanation*, to give an account of how and why the distinctive behaviour of generic claims arises. That there is a connection between generics and default inference is a relatively standard view [24, 32, 47], although there is less agreement about precisely how that connection is to be understood. The distinctive contribution of this paper is applying this connection to the social phenomenon of the formation of the common ground in discourse. This will help

*Email: gr69@st-andrews.ac.uk. Thanks to colleagues, including Deb Brown, Chris Cordner, John Flett, Kai von Fintel, Rohan French, Anil Gupta, Sally Haslanger, Lloyd Humberstone, Karen Jones, Bernhard Nickel, Graham Priest, Dave Ripley, François Schroeter, Laura Schroeter, Mandy Simons and Shawn Standefer, and audiences at Melbourne, the Graduate Center at CUNY, Pittsburgh, UConn, MIT and UQ for helpful feedback as I worked through this material. ¶ This research is supported by the Australian Research Council, through Grant DP180103687.

us account for how we come to accept characterising generics, even when they are not the explicit topic of discussion. We accept generics by a well-understood phenomenon of discourse accommodation [2, 41, 43], applied to *inference*. This connection, between accommodation, inference and characterising generics, will give us some better tools to engage in *improving* our use of generic claims, and what is at stake when we argue about them.

1. CHARACTERISING GENERICS

Generic claims have the form *Fs are G*, where there is no explicit quantifier expression. So, claims like *some men are violent*, and *all men are violent* are not generic claims, and neither are claims like *most men are violent* or *normal men are violent*. In this chapter, I will focus my attention on bare generic claims of the form *Fs are G*, though there is a closely related form with definite or indefinite descriptions *the F is a G* (e.g. *the bird flies*) or *an F is a G* (*a bird flies*).

Our focus will be on so-called *characterising* generics, in which items of class *F* are characterised as generally having feature *G*. Each of the examples we have seen so far have this function, but not all claims of the form *Fs are G* characterise individual *Fs* in this way. Consider *birds are widespread*, and *children ate all the pizza*. These claims do not characterise individual mosquitos as widespread, and no individual child is said to have eaten all the pizza. These *class* generics describe the behaviour or features of a group as a whole, rather than characterise their members individually. It will be enough for us to be going on with to focus on characterising generics, because these play a significant role in our thought and talk (especially when it comes to social kinds), and the most difficult issues concerning the truth conditions of generic claims concern characterising generics.

So, to clear the ground, when we consider characterising generics, such as

- *Birds lay eggs*
- *Mosquitos carry RRF*
- *Logic talks are boring*
- *Cows are food*
- *Men are violent*
- *Tall people with back injuries find it difficult to drive small cars*

each of which have the form *Fs are G*, and each characterises members of the class *F* as having feature *G*. In each case, to make the generic claim you do *not* thereby say that all *Fs are G*. (It is true that birds lay eggs, but male birds do not.) You do not merely say that some *Fs are G*. (*Some men are gentle and are not violent*. This does not necessarily justify the generic claim *men are gentle*.) The generic claim *Fs are G* does not even require that most *Fs are G*, or that normal *Fs are G*. (Most mosquitos do not, in fact carry RRF. Male mosquitos do not, and neither do mosquitos outside Oceania. These mosquitos are, nonetheless, normal.) Finally, while it might be correct to say that mosquitos are the *kind* of thing that carries RRF, this will only help

in specifying the truth conditions of characterising generics when supplemented with an account of kinds and their properties.

As the last example in our list indicates, we can form characterising generic claims with complex predicates. It is one thing to think that semantic and cognitive competence requires some sort of commitment to a metaphysics of kinds such as *birds*, *men*, or *logic talks*. It is another for that account to extend to complex kind terms such as *tall people with back injuries*, or *Dundee residents who are sympathetic to independence but who weren't included in this survey*.

It is important to underline some of the distinctive behaviour of characterising generics. First, they do seem to be truth apt. It is not only meaningful to assert *birds lay eggs*, it is also meaningful to make claims like these –

If *birds lay eggs* then birds and reptiles share a feature.

If that talk was representative, then *logic talks are boring*.

–so we can meaningfully use characterising generics in the antecedent or in the consequent of conditionals.¹ It is certainly possible for any generic *Fs are G* to form the *opposite* generic, *Fs aren't G*, though it is less clear that we have a natural way to assert the *negation* of a characterising generic, other than the stilted phrasing *it's not the case that Fs are G*. Consider the characterising generic *birds are female*. This seems untrue (to me, at least), when understood as a characterising generic, since birds are sexually dimorphic. The opposite claim *birds are not female* (when also understood as a characterising generic) seems equally untrue, and this seems good evidence that *Fs aren't G* does not do a good job of expressing the negation of *Fs are G*. The ungainly wording *it's not the case that Fs are G* (or something like it) seems to be required if we wish to state the negation of the characterising generic.

The 'logic' of characterising generics, such as it is, is complex. We have already seen that *Fs are G* does not mean that *all Fs are G*, or that *most Fs are G*, or even that *normal Fs are G*. As one example of the strange behaviour of characterising generics, even though *birds lay eggs* is true, and *all birds that lay eggs are female* is true, it does not follow that *birds are female*. In other words, from *Fs are G*, and *All Fs that are G are H*, it need not follow that *Fs are H*.² In characterising generics, we cannot, in general, *weaken* the consequent.³

¹Unlike imperatives, which cannot be used as the antecedent of a conditional, as one example. It makes sense to ask "pass the milk" but no sense to say "if pass the milk then . . ."

²This means that the characterising generic acts significantly different from any *any* account which posits the truth conditions of *Fs are G* in terms of the existence of some selected sub-class of *Fs* (whether *most Fs*, *normal Fs*, *contextually chosen Fs* – where the class is chosen independently of the selection of the predicate *G*), such that all *those Fs are G*, for if *all Gs are H*, then any such selected collection of *Fs* that are each *Gs* will also be *H*. Any treatment of *Fs are G* by way of a contextually restricted quantifier will need to allow for the characterising predicate *G* to help determine the selected restriction.

³This makes generics even stranger than counterfactual conditionals. If a counterfactual conditional $A > B$ is true (that is, if *A* were the case, then *B* would be the case), and if *B* entails *C* (that is, in *any* possible scenario in which *B* holds, so does *C*), then it follows that $A > C$ is true too (that is, if *A* were the case, then *C* would be the case). The analogous rule fails for characterising generics.

Since the “logic” of generics seems wild, any reasonable account of their truth conditions turns out to be complex. Here is one example from the recent literature, in a paper by Sarah-Jane Leslie [23, page 43]:

Though there may be a further refinement or two needed, we can describe the circumstances under which a generic of the form ‘Ks are F’ is true as follows:

The counterinstances are negative, and:

If F lies along a characteristic dimension for the Ks, then some Ks are F, unless K is an artifact or social kind, in which case F is the function or purpose of the kind K;

If F is striking, then some Ks are F and the others are disposed to be F;

Otherwise, almost all Ks are F.

I would suggest that these worldly truth specifications—these descriptions of how the world must be for the sentence to be true—should not be mistaken for semantically derived truth conditions, however.

The complexity of these truth conditions are due to the wild behaviour of generic claims. The reasons for each particular component need not detain us here. However, we should note that characterising generics with complex kind terms like *tall people with back injuries* and *Dundee residents who are sympathetic to independence but who weren't included in this survey*, put pressure on the division into artifactual, social and natural kinds, and as Leslie hints, her account may need to be refined to deal well with generic expressions with complex kind terms. I think it follows from this that the connection between the use of generics and the propensity to *essentialise* categories is quite subtle, and worth further reflection [17, 18, 25].⁴

As Leslie points out, these are “worldly truth specifications” – characterisations of what it might be for the world to satisfy a characterising generic of the claim Ks are F – and are not to be taken to be *semantically derived* truth conditions, which in some sense mirror the structure of the ground-level claim. Leslie uses the following example to illustrate the point. The claim that *Bob is red* may be true if and only if Bob is experienced as red when observed by standard observers in standard conditions. This is a statement of worldly truth conditions for the claim. However, the *semantically derived* truth conditions will not have this form, because semantic competence with the claim *Bob is red* need not involve any commitment concerning standard observers or standard conditions. The semantically derived truth conditions may have the form of an object (in this case, the referent of *Bob*) falling under a the extension of a predicate (in this case, the extension of *red*). Leslie argues that the *semantically derived* truth conditions for characterising generic

⁴I note here that the analysis of generics given here is completely orthogonal to the question of whether generics in any way *essentialise*. The question will become: whether and when – and how – do inference and explanation *essentialise*?

claims may simply *use* the generic quantifier. When it is true that *Fs are G*? If and only if, $(\text{Gen } x)(Fx, Gx)$, where ‘Gen’ is the generic quantifier, used in the meta-language.

2. METASEMANTICS

In one sense, such deflationary truth conditions for generics *have* to be correct (at least, they have to be, if characterising generics *have* truth conditions), but they do not help when it comes to explaining competence with generic expressions, or what capacities one has to employ in order to be able to use characterising generics. If the truth conditions for generics are either too complicated to be able to account for competence (as Leslie’s example worldly truth conditions are), or too *thin* to be informative (as the deflationary truth conditions are), perhaps we should look elsewhere.

Inferentialism is an approach to semantics that attempts to give an account of the semantics of expressions in terms of *norms of inference* [6, 7]. Inferentialism, in this sense, is a species of the wider genus of *Normative Pragmatics*, which encompass approaches to semantics that centre on *norms of use* (perhaps including inference, perhaps not) [19, 21, 27]. A normative pragmatic account of the semantics of some concept will be a different way of characterising its meaning, in contrast to a truth-conditional semantics. For a straightforward example, an inferentialist account of the semantics of *conjunction* may characterise conjunction as follows: from the conjunction of A with B you can infer A and you can infer B. Conversely, from A and B together, you can infer the conjunction of A and B. This is an inferentialist account of the semantics of (sentential) conjunction. A truth-conditional account of the semantics of conjunction goes as follows: the conjunction of A and B is true if and only if A is true and B is true.

It should be clear from this toy example that an inferentialist account of the semantics of conjunction and a truth-conditional account of the semantics need not be in conflict. One might endorse both as correct [36, 38]. It is a further question which, if any, of these two kinds of accounts might have priority in any order given order of explanation. However, since the truth conditions of characterising generics seem ill-suited to explain our competence with generics,⁵ and since generics are something we can *communicate* with and *coordinate* on, there is scope for looking elsewhere for an explanatory rich framework for understanding generics, and what is at stake when we disagree about them.

To explain further the difference between kinds of semantic explanations, consider the concept of *possibility* and *necessity*. *Possibly p* and *Necessarily p* have truth conditions expressed in terms of *possible worlds* of the familiar shape: *Possibly p* is true (at some given possible world *w*) if and only if *p* is true at some possible world *v* (perhaps *v* must be possible ‘relative to’ *w* in some sense). *Necessarily p* is true

⁵In particular, we should attempt to explain the fact that we acquire generics very early in our cognitive lives, plausibly significantly earlier than we acquire the ‘simpler’ logically precise universal and existential quantifiers [23, p. 21], which have much simpler truth conditions.

(at some given possible world w) if and only if p is true at *every* possible world v (perhaps, again, these worlds v must be possible ‘relative to’ w). These are the standard *truth conditions* for the concepts of possibility and necessity. However, we do not acquire the concepts of possibility and necessity by first acquiring the concept of a possible world and the notion of truth ‘at’ a world. Arthur Prior puts the point this way:

To say that a state of affairs obtains is just to say that something is the case; to say that something is a possible state of affairs is just to say that something could be the case; and to say that something is the case ‘in’ a possible state of affairs is just to say that the thing in question would necessarily be the case if that state of affairs obtained, i.e. if something else were the case. . . . We understand ‘truth in states of affairs’ because we understand ‘necessarily’; not *vice versa*. [35]

According to Prior, we *first* understand the concepts of possibility and necessity, and then we use these to understand what it takes to be a possible world (or, in his vocabulary, a possible state of affairs). This leaves open the question of how we understand the concepts of possibility and necessity, and in particular, why it is that the concepts we have end up having the truth conditions they in fact have. I have argued elsewhere that an inferentialist account of these concepts gives us a plausible account of *how* we could acquire them, and use them in such a way that they end up having the expected truth conditions, *without* requiring that we first grasp the concept of a possible world [37, 39]. The inferentialist semantics for the modal concepts exploit our ability to not only assert and deny claims categorically, but to assert and deny (and infer) *under the scope of a supposition*, including subjunctive suppositions (where we suppose things had gone differently to how things actually went), or indicative suppositions (where I might suppose, contrary to my current view, that you are in fact, correct). Once we notice that our practices of thinking and talking can run “offline” under the scope of such context shifts, they have the structure that allows for us to introduce modal expressions.⁶ The particular details are not important, the everyday context shifts in discourse (and cognition) given by making claims under the scope of different kinds of supposition, allow for the introduction of modal concepts in such a way that we can see how we can acquire modal concepts in our own *thinking*, and can coordinate on them in our *conversation*, without requiring that we have any antecedent understanding of what it takes to be a possible world, or what it takes to be true *in* such a thing.⁷

⁶As one example, to show that p is *necessary*, you show p in an ‘arbitrary’ context, that is, in a fresh context in which no other assumptions are made. I have given more details about how inference rules like these determine the behaviour of these modal concepts elsewhere [37, 39].

⁷There are, of course, connections between conversational or cognitive *contexts* and possible worlds, but they are not the same. For one thing, when we consider what it would be like if kangaroos had no tails, we might suppose that kangaroos had no tails, and start thinking or conversing under the scope of that supposition. That is a fresh context. We are not thereby describing a different possible world, on the usual understanding of worlds, for there are *many* different possible worlds in which kangaroos have no tails, or, if it is necessary that kangaroos have tails, then there is

* * *

An inferentialist account of the significance of characterising generics takes a similar approach: rather than first attempting to give the truth conditions of a characterising generic statement, we explain its semantics in terms of inference. This seems like a natural thought, for it is widely recognised that there is some kind of connection between generics and default, defeasible inference [24, 32, for example] and default, defeasible inference itself plays a central role in inferentialist semantics [6, 7], so the connection seems plausible. At the very least, it is natural to think that while the inference from $Fs \text{ are } G$ and Fa to Ga is not deductively *valid*, the inference is appropriate as a *default*. This seems to be the natural way to characterise the inferential power of making the claim $Fs \text{ are } G$. When it is appropriate, though, to make such a claim? The natural inferentialist answer is the converse: when we can make the default inference from Fa to Ga , when you make no particular assumptions about a . This is the simple and direct inferentialist account of the meaning of characterising generics, and it has the same form as the inferentialist account of necessity and possibility, mentioned above.⁸ Though this account is relatively natural, it has faced some tough criticism, especially from those, like Pelletier and Asher, who take the connection between generics and default inference seriously.

. . . default logic does not provide us with an acceptable formalization of generic statements. Default rules are rules, and therefore are sound or unsound – rather than sentences, which are either true or false. If we analyze characterizing sentences using default rules, these sentences would not have truth values, and their meanings could not be specified by an ordinary semantic interpretation function. One consequence of being neither true nor false is that characterizing sentences would therefore not talk about the world, instead they would talk about which inferences to draw. And this seems to us to be a strike against such an account. [32, p. 1152]

It should be clear, given the example of the inferential account of modal operators, that this criticism is misplaced if used to target *this* inferentialist account of characterising generics. The idea is not that a generic claim $Fs \text{ are } G$ is true if and only if we can make the default inference from Fa to Ga . That would be to use the inference rule as a truth condition, but the inference rule does not attempt to give truth conditions. The inference rules tell us when we can infer something *from* a characterising generic, and when we can *conclude* a characterising generic. Like

no possible world in which kangaroos have no tails. In any case, there is a single fresh conversational or cognitive context introduced when we suppose that kangaroos had no tails, and these contexts are not in a one-to-one relationship with possible worlds.

⁸It also has the same form as the inferentialist account of the universal quantifier. What can you conclude from $\forall x\phi(x)$? You can prove $\phi(t)$, for any singular term t . When are you in a position to prove $\forall x\phi(x)$? When you can prove $\phi(a)$ for some arbitrary term a .

the other sentences of the language, characterising generics are truth apt simply because they are, on this account, the kind of things that can feature as premises and conclusions in reasoning. The inference rules, however, do not take the form of truth conditions.⁹ It is not my place to give an extended defence of the inferentialist semantics for characterising generics. Thankfully, that work has been done by Preston Stovall in his “Characterizing Generics are Material Inference Tickets: a proof-theoretic analysis” [47]. In the remaining sections of this chapter, I will take the broad brushstrokes of this view as given, and show how the connection between default inference, the common ground and accommodation phenomena can help explain the difficulties that arise around disagreement involving characterising generics, and how we might, nonetheless, make progress.

3. WHAT IS INFERENCE?

Since inference plays a central role in any inferentialist semantic explanation, it will be useful to spend some time clarifying what inference involves. In particular, it will be *very* important to distinguish the material notion of *default* inference from any *formal, deductive logic* understanding of the term. To get an initial sense of the significance of inference, consider the difference between (1) and (2) below:

- (1) Tweety is a bird. Tweety flies.
- (2) Tweety is a bird. So, Tweety flies.

In both cases, the two claims are made, that Tweety is a bird, and that Tweety flies. But in the second case, the second claim is *inferred* from the first. Here, the additional component is that the first claim, to the effect that Tweety is a bird, is presented as *justifying* the second claim. An inference, here, is an action, in which a claim (here, the conclusion, to the effect that Tweety flies) is “backed up” by putative evidence (here, the premise, to the effect that Tweety is a bird). In this example, the premise came before the conclusion, but there is no need for the two components to be presented in that order. We could just as easily have said:

- (3) Tweety flies, since Tweety is a bird.

thereby presenting the fact (if indeed, it is a fact) that Tweety is a bird as justification for the claim that Tweety flies. As far as public speech, dialogue, or written text goes, inference can be presented in either direction. There is no requirement that the conclusion *conclude* the inference. It can be stated up-front.

⁹The distinction is most stark in the case of the natural deduction inference rules for the logical quantifiers. These rules *clearly* characterise the semantics of the quantifiers \forall and \exists . They do not do so by taking the form of truth conditions, because they do not involve any characterisation of a domain of quantification, or make use of assignments of values to variables, or any of the paraphernalia of models for first-order predicate logic which are the appropriate way to characterise the quantifiers truth-conditionally. Despite that, the inference rules do manage to inferentially characterise the behaviour of the quantifiers in a way that agrees with the truth-conditional account. The two accounts differ, but are not in any conflict.

What goes for public speech can also go for private thought. In my own reflection, I can come to think that Tweety is a bird, and conclude *from this* that Tweety flies. Or, I could come to believe that Tweety flies, and then wonder if that is really the case, and reassure myself about this fact, since I also know that Tweety is a bird. We can equally understand inference as a kind of transition in our *talk*, and as a transition in our *thought*. In what follows, I will not take a stand as to which of these options, if any, is the fundamental or original notion of inference. Everything in this account of generics will be consistent with the view that inference in thought is *internalised* public inference, and that the norms governing inference are fundamentally a social, communicative matter.¹⁰ And, everything in this account is consistent with the view that we can think of inference as presented in language as *externalised* representation of a more fundamental notion of inference in thought.¹¹ Or, you could think that this distinction is somehow ill-posed, and that neither a thought-first nor a talk-first account is correct. We need not take a stand here, and I draw attention to it merely to clarify that the ‘inferentialism’ important to this account is a thin one, according to which inference (whether understood thought-first or talk-first) can play a role in giving an account of the meanings of characterising generics.

To make out this case, it is important to distinguish our target notion of *inference* from *proof*, *logic* or *deductive reasoning*. The inferences recorded in (1), (2) and (3) above are not *proofs*, and nor need they be treated like proofs by the one who infers in that way. They are not logically valid, and they need not count as deductive reasoning. The inference recorded here is *invalid* in the formal logical sense, because it is formally possible that ‘Tweety is a bird’ be true while ‘Tweety flies’ is false.¹² Even if we impose some realistic semantic constraint on the interpretation of the predicates ‘flies’ and ‘is a bird’ the mere possibility of non-flying birds (emus, hatchlings, injured birds, etc.) ensures that counterexamples to the inference abound. Nonetheless, inferences like these – in which we make everyday transitions from premises to conclusions, despite the possibility of counterexample – are made everyday, and are a part of competently using concepts such as *bird* and *flies*. The relevant notion inference for our account is *contingent* and *material*, and neither *necessary* nor *formal*.

Closely allied to the material nature of inference, is its *defeasibility*. Take the inference from ‘Tweety is a bird’ to ‘Tweety flies’. In a given context where I (not knowing what kind of creature Tweety happens to be), make a justification request for your claim that Tweety flies. You say, ‘Tweety is a bird’ to discharge that request, thus making that inference. I am well within my rights to take that justification request to be met, and to carry on with the conversation. If you did not finish

¹⁰This is the normative pragmatic inferentialism of Robert Brandom [7].

¹¹Paul Boghossian’s account of inference [5], as one example, takes it that inferential transitions in thought are fundamental. Similarly, Gilbert Harman’s influential account of reasoning in *Change in View* [5] is all about the dynamics of change in *beliefs*.

¹²The argument has the first-order logical form Fa therefore Ga , so it is manifestly formally invalid.

by saying that Tweety is a bird, but added ‘. . . in fact, she is an emu’, I would be well within my rights (knowing, as I do, that emus do not fly), to take it that my justification request is now no longer met. What counts as a good inference in some context may be defeated in the presence of new information.

This should not surprise us, if we think of the making of inferences as a way to impose some degree of quality control our thought and our talk. There is reason to take up claims from others and from own thoughts and to commit to them for ourselves (it is hard to see how we could get going in our own cognitive lives were we to do otherwise) and there just as much reason to attempt to ask some questions concerning claims that seem, well, *questionable*. Any process like this will look something like the making of justification requests, in the way I have sketched it out here.¹³ It is not at all surprising, given the role of quality control, that what counts as meeting a justification request for a questionable claim might vary from context to context, and that inferences may be defeasible in the way that we have seen. After all, sometimes a proffered explanation can go too far and raise more questions than it settles, such as when you go on to explain that Tweety is an emu, or that she is a bird with a wing injury. None of this means that the original explanation (that she is a bird) would not have done enough to assuage us. In wanting our questions answered, we do not necessarily need *every* question to be answered. We can treat a claim as justified without meeting the logicians’ standard of exceptionless deductive validity.

Nonetheless, we can see why the notion of inference can lead to the logicians’ sense of validity as a kind of limiting case [9–11, 15], since there are two different ways that the answer to a justification request can fail. Your answer to my request for a justification can be rejected as a claim that is itself *untrue*, or itself *unjustified* (say, I ask you to justify your claim that Tweety is a bird). However, I can grant that your further claim is in order, but reject it as not meeting the claim to justify the conclusion. I could be the skeptic and say, I grant that Tweety is a bird, but isn’t it possible that nonetheless, Tweety might not fly? The logicians’ sense of *deductive* validity is an attempt to mark out a kind of limiting case of good inference, where the offered justification leaves *nothing* out, and anyone who accepts the premises but still takes the conclusion to be in question has shown that there is some kind of failure to communicate, rather than a disagreement whether to say ‘yes’ or ‘no’ to a shared question.¹⁴ For our purposes, however, we need not concern ourselves with this idiosyncratic case of deductive logic.

Before turning to the connection between generics and inference, let’s pause to notice the breadth of uses we find for our inferential capacities. The ‘quality control’ nature of making justification requests can be applied, usefully, in other domains. An important one for creatures like us is in our *planning* and *decision-*

¹³For an account of dialogue which takes justification requests seriously, seeing it as a kind of quality control over the commitments made by each interlocutor, see Charles Habmlin’s *Fallacies* [16, Chapter 8].

¹⁴If I come across someone who grants *A* together with *if A then B* but still thinks that *B* is in question, then it seems that they must use “*if . . . then . . .*” in some way that I do not yet understand [8].

making. In planning for the future, we consider different ‘possibilities’. This need not be understood as a kind of observation, where we come across a thing called a possibility and attempt to describe it as best we can. Instead, we can understand planning as inference from suppositions. When we consider what we take to be a future possibility, we *suppose* it comes about, and infer from there, as best we can. We apply our practice of inference, drawing out consequences, not to something we take to be the case, and not to something someone else has asserted. Instead, we apply the same reasoning processes, asking ourselves what else would follow *were* that possibility to come about. In this way we can do some thinking in advance, and make decisions now, to apply in future circumstances, if they arise, or use our understanding of the significance of future possibilities to inform our decisions about options to take [20]. Our inferential capacities are at work in each of these ways of engaging in with the world, and so, given the connection between characterising generics and inference, are found at the centre of these modes of engagement.

4. GENERICS MAKE DEFAULT INFERENCE EXPLICIT

So, with this understanding of inference and its role in discourse and thought in view, we can more clearly understand the connection between default material inference and characterising generics. The key thought is that default inference is acquisitionally and conceptually prior to our characterising generics. The force of a characterising generic *Fs are G* is to license the default inference from *Fa* to *Ga*. If I grant *birds fly*, I ought not demur at the default inference from *Tweety is a bird* to *Tweety flies*. Conversely, if I accept the default inference from *Fa* to *Ga* (where I appeal to no particular features of *a* other than the assumption that *Fa*), then I am in a position to grant the characterising generic *Fs are G*. This does not mean that the characterising generic is *about* whatever inferences we accept, since these statements about the norms for granting generics, how we are in a position to infer to a generic, and what we can infer from them, and are not statements of their truth conditions.¹⁵ It should not come as a surprise that an account in terms of inferential power should itself involve the inferences that we license.

The power of this analysis is that it explains the key distinctive behaviour of characterising generics. Characterising generics have exceptions (birds fly, but penguins do not, despite being birds), in just the same way that default inference has exceptions. We have an explanation of why default inference comes with exceptions: it would be a waste of time and energy to demand that all justification requests be met only at the highest, 100% standard. We count justification requests as met with a much lower threshold, accepting reasons that we take to count in *this*

¹⁵See the earlier discussion of the natural deduction rules for the logical quantifiers. These rules state that to *prove* $\forall x\phi(x)$ from some set of assumptions, I can prove $\phi(a)$ for an *arbitrary* *a* (a name ‘*a*’ that does not appear in the assumptions used). This rule characterises the universal quantifier, and gives it a meaning that gives it the usual truth conditions. So, it need not follow that these inference rules give the universal quantifier deviant meaning that is somehow ‘about’ arbitrary names, or provability from premises.

case, even though they might not count in *every* case. Since granting an inference is taking a justification request (whether implicit or explicit) to be met, then we have an explanation of why it is that default inference carries exceptions. Since generics inherit features from default inference, we can see why it is to be expected that characterising generics have exceptions, too.

Let's try this for another example of strange behaviour of generics: the failure of weakening of the consequent, mentioned earlier, when we introduced generics. We grant *birds lay eggs* and we grant, also, that *all egg-laying birds are female*. It does not follow, though that *birds are female*. We resist granting this characterising generic, despite acknowledging that all egg-laying birds are female. Why do we resist? This is easily explained by appealing to how we treat justification requests and inferences. The inference from *Tweety is a bird* to *Tweety lays eggs* is acceptable, at least in some wide range of contexts. If I am unsure of what kind of animal Tweety is, and I am surprised that she has laid an egg, you can explain that she's a bird, and my question is answered. If the relevant options are that Tweety is a bird, or a mammal, or a reptile or some other creature, then settling on *bird* helps settle the issue between egg-laying and live birth. It is *much* harder to find a context in which making the inference from *Tweety is a bird* to *Tweety is female* is anywhere near as helpful. In most contexts in our thought or our conversation, we are not going to meet a justification request for the claim that Tweety is female by granting or learning that she is a bird.¹⁶

Let's consider the phenomenon of low-rate generics with striking properties. Despite the fact that very few mosquitos carry Ross River Fever, it makes sense for us to grant *mosquitos carry RRF*. One explanation may be that these are not, at root, characterising generics in which we characterise individual mosquitos, but a class generic, where we ascribe a property to the population of mosquitos as a whole. This population carries RRF, in the same way that this population is *widespread*. However, the cases are also dissimilar, in that individual mosquitos do carry RRF, so let's see whether an inferential treatment of these cases can make sense of this phenomenon.

While it would odd to make the inference from a given claim *x is a mosquito* to *x carries RRF* in any given context in which a mosquito is under consideration, there are important contexts in which the inference does make sense, specifically when we are asking questions concerning how it is that someone caught RRF. We can answer a justification request for "I caught RRF" by "I was bitten by mosquitos", and this clearly makes sense in many conversational contexts. Similarly, if we are con-

¹⁶This is not to say that you cannot cook up artificial contexts in which this inference makes sense, of course. Take a background in which we are considering the animals in some large factory farming facility. There are humans, birds, and maybe some domestic dogs. The birds are all battery farmed chickens, and these are all (or are mostly) female, caged for their eggs. In this context, where I consider a randomly chosen creature in this facility, it makes sense for me to make the inference from *x is a bird* to *x is female*. In just the same way, it is appropriate *in these contexts* to grant *birds are female*, where this is understood as appropriately contextually restricted, to when we are talking about birds *around here*. The fact that generics seem *as appropriate* as the corresponding inferences in contexts such as these seems to be evidence in favour of the generic/inference connection.

cerned about RRF, and wish to know how it is transmitted, it makes sense to be told that it is *mosquitos* that carry RRF, rather than flies or midges. The explanatory and justificatory connection makes sense, even when the rate at which the property is held in the population is quite low, especially in cases like this where we are keen to protect ourselves from danger.¹⁷ Given the importance of making the inference in cases such as these, we can see why it can make sense to grant the generic *mosquitos carry RRF*.

It is worth reflecting, before we end this section, on why it is that characterising generics have the form *Fs are G*, if they have the role of making inferential transitions explicit. After all, in our inference, we can make a transition from any claim to any other claim – given any two thoughts, there may be some context in which the transition from one to the other makes sense. If we were to make an inference from A to B explicit, the appropriate construction would be a more general *conditional* construction which would allow for arbitrary propositions as antecedent and consequent, of the shape *if A then B*. Why do we gravitate toward the generic *birds fly*, which take predicate expressions, rather than more general conditional constructions, which take sentences as components? We don't say *if it's a bird, it flies* with anything like the frequency of *birds fly*. It is a lacuna in this account of generics if it can provide no answer to this puzzle.

The first component of an answer is that generics of the form *Fs are G* express rules of thumb that are *about* Fs, as a means of cataloguing an expressing our understanding of Fs and our orientation towards them. Given our interest in responding to the world around us and the use of our schemes of conceptualisation of our environment, it is not surprising that we have basic means of representation that hews closely to the subject–predicate form. In our thought and talk about birds, when we are making inferences about them, the subject matter is fixed. The transition between being a bird and flying can apply to Tweety, or to other birds. We do not change the subject or item described when we make the transition from premise to conclusion (or conclusion to premise). Given that much of our inquiry involves figuring out how things are, the subject matter often does not change from claim to claim, so it is no surprise characterising generics *Fs are G*, in which the *selected* item (the F) is also the item characterised (as a G) have a central role in our conceptual architecture.

One important caveat to keep in mind is that nonetheless, is that our thought and talk, at a very basic level, *does* involve means to express broader inferential connections in which the selected target (whatever is picked out 'antecedent') is the thing described (the circumstances of the 'consequent'). Deeply held conventional wisdom of the form *red sky at morning, sailors take warning* is but one candidate for a claim we use to express a generic-like default conditional connection, in which the antecedent describes one thing (the colour of the sky in the morning) and where the consequent describes another (the activity of sailors). This has many of the

¹⁷See Sarah-Jane Leslie's "The Original Sin of Cognition" [25] for more on the phenomenon of low-rate striking property generics.

features of characterising generics. If on the occasional red-skied morning, sailors recklessly go out on the ocean, failing to heed the ominous signs, this need not be taken as a refutation of the rule of thumb.

So, with the widespread use of characterising generics in our cognitive and communicative practices, together with allied phenomena which can equally be understood as playing an expressive role to make explicit in our thinking and our conversation the kinds of inferential transitions we make in our thought and talk.

Another striking feature of generics is the difficulty we have in expressing a denial of a generic claim. Consider, for contrast, the universal quantifier. To deny the claim *all Fs are G* it suffices to assert *not all Fs are G*. Consider a comparable assertion that expresses a denial of the generic *Fs are G*. This is not done by way of the opposite generic *Fs aren't G*, at least if that has the meaning that, generically, Fs fail to have the property G. After all we can reject both *Fs are G* and *Fs are non-G* when rejecting both inferences, from something being an F to its being a G and to its not being a G. One can reject the generic *birds are female* and *birds aren't female* equally. How can we express this rejection by way of asserting a negative statement? The clearest way to do so is something like *it is not the case that birds are female*, where the negation clearly takes wide scope over the generic statement, but even this can be plausibly understood as an inner negation, claiming that *birds aren't female*. Given the difficulty of fixing scope with generic expressions, it can be much clearer to make some quantifier expression explicit, to say *many birds aren't female* to indicate why we resist the generic *birds are female*, but as we have seen in the case of low-rate striking generics such as *mosquitos carry RRF*, or in the response *not all men* to accusations to the effect that men are rapists, holding to generic can resist the presence of an explicitly quantified negative claim in response, so this strategy cannot be relied upon to provide a clear statement of an opposing position. This striking behaviour of negated generics makes disputes in which generics are the subject matter particularly intractable.

In the following section, we will see that this issue becomes all the more pressing when we see that we do not just *learn* generics when we come to consider them for ourselves and decide whether we agree with them or not. Instead, we can find ourselves committed to generic claims by a kind of conversational osmosis from our surrounding community. It is to this phenomenon that we will now turn.

5. ACCOMMODATION & INFERENCE

Conversation is a shared exercise, in which the participants build something together. A part of what they build is the *common ground* of that conversation, those things that everyone in the conversation has granted, and from then on, can take for granted. The notion of common ground has been central in the contemporary study of the interaction between semantics and pragmatics in recent decades [14, 18, 43, 45]. Kai von Stechow describes the common ground as follows:

The COMMON GROUND of a conversation at any given time is the set of propositions that the participants in that conversation at that time

mutually assume to be taken for granted and not subject to (further) discussion . . .

When uttered assertively, sentences are meant to update the common ground. If a sentence is accepted by the participants, the proposition it expresses is added to the common ground. [12]

What is shared between the participants of a discourse includes more than a body of propositions. Since conversations involve questions, we also keep track of a shared list of questions currently under discussion [4, 40, 42], and if the conversation includes imperatives, we will also together keep track to-do list for each individual [33, 34]. Furthermore, the fact that conversation is a shared enterprise does not mean that each participant's own private commitments are transparent to the others. There is a difference between the public common ground (what we have committed to in public, by way of what we have said and what we have let stand) and our own private commitments [28]. A public record of a common ground is not inconsistent with deceit (where we assert something in conflict with what we hold true, for the purposes of manipulating others) and neither is it inconsistent with silent dissent, in which we let what someone else says stand, without raising an objection, even though we privately disagree. In both cases, what is *said* in the conversation can stand as the common ground of that discussion, despite this not being a record of the private commitments of the participants.

One important feature of the common ground, in understanding the dynamics of discourse, the phenomenon of *accommodation*. The shared propositional state of the conversation – those claims that we together take for granted – can shift not just by updating with the content of what any participant explicitly asserts. Other items can enter into the common ground in order to accommodate what has been said. If, after giving a presentation as a visiting scholar at a research seminar, I say *if there is dinner after the seminar, my son will come*, then, if that statement meets with no objection, the common ground of that conversation will be updated, not only with the conditional claim, but also with the information that I have a son, and most likely, that I have *one* son. I did not literally say that I have a son, let alone that I have only one. Speaking literally (if unhelpfully), the content of what I said could be true even if I had been childless, had it been certain that there were no dinner after the seminar, and so the issue of anyone coming with me would not arise. Of course, no-one would interpret such a statement in that way, and we all take it for granted that I have the son that I was mentioning in the claim. If my claim passes into the common ground with no objection, then another participant in the conversation can meaningfully ask *how old is he?*, and it is obvious to all concerned that the question is about my son, who we have all heard about in my previous claim. This phenomenon is called *presupposition accommodation*. The details of *how* – and *which* – presuppositions are accommodated is a matter of active research [3, 13, 41, 44], and we need not go into the details of any account of presupposition accommodation here. Suffice it to say that it is now very much a live issue to explore the dynamics of conversation and the way that our commitments are updated not only by way of

the content of what is explicitly asserted but any number of ways that are less explicit. This area of research has included rich connections to matters of ethics and normativity more broadly construed, too [1, 22].

Given the connection between characterising generics and inference I have delineated above, we can see that generics, too, are added to the common ground in conversation, in ways that are reminiscent of accommodation phenomena, but have their own distinctive features. The point is a simple one. On the standard picture, if I say *Tweety is a bird. So, she flies*, in conversation, and I meet with no objection, then the common ground is updated with the information to the effect that Tweety is a bird, and that she flies. However, my statement did more than just make the assertion that Tweety is a bird, and the assertion that she flies. In addition, I treated the connection between these claims as an inference, and this, too, is something that *could* meet with resistance, so the inference could, in its own rights be rejected, even without rejecting the two assertions. It makes sense, then, to think of the inference as being recorded, somehow, in the common ground, to mark the difference in our conversational commitments in the case where an inference is rejected from a case in which it is not. We have a candidate for how the common ground can be updated in cases where we make the inference from Fa to Ga and meet with no resistance: it is that the common ground is updated with the generic $Fs \text{ are } G$. At least in this local conversational context, if we grant the inference from Fa to Ga (so, in this conversation we appeal only to Fa to answer the justification request for claim Ga), then we have granted, implicitly, the generic $Fs \text{ are } G$. If, in a conversational context, we can appeal to Fa alone as a justification for Ga , then at least relative to this context, we are committed to the corresponding generic.

If this view is correct, then the common ground updates with the generic $Fs \text{ are } G$ whenever an inference from Fa to Ga is admitted without objection. This is like presupposition accommodation in the way that a propositional content enters the common ground without being directly asserted. However, the propositional content that is admitted (the characterising generic) is not some presupposition of something else asserted. It is the propositional content that corresponds to the making of inference itself. The inference from Fa to Ga is licit, in this context, if (in this context) $Fs \text{ are } G$. The proposition to update is directly represented in the conversation by the inference (the 'so', 'therefore' or however else the inference was represented), and there is no complicated negotiation to expand the common ground appropriately to meet the presuppositions of some other asserted content. The common ground updates in a manner that is explicitly represented in the discourse itself.

However, although the points at which inferences are made in discourse are explicit, and can be questioned and rejected in just the same manner in which assertions are refused admission into the common ground, blocking an *inference*, in particular, is not a straightforward matter. If I say Fa , *so* Ga , it is easy enough for you to block either of the assertions of Fa and Ga on their own: you simply call either assertion into question with an 'are you sure?' or 'I don't agree'. It takes more

finesse to single out the *inference* as faulty, especially when you actually grant Fa and Ga on their own terms, but you take it that Fa doesn't count in favour of Ga. Take an example circumstance where someone from a minority (say, a logician, in the Philosophy Department) has given a boring talk, and a conversation partner says 'of course it was boring, it was a *logic* talk', then in the context where the talk did happen to be boring, and it is not in question that it was, indeed, a logic talk, to object to the inference you have to make the connection more explicit than has been said in the inferential transition from one claim to the other, by saying something like 'the fact that it was boring has nothing to do with it being a *logic* talk.' Simply saying 'no' to the inference does not have the required precision, because the refusal could equally be interpreted as a refusal to accept either of the assertions instead of as an objection to the inference itself.

Inferential transitions go by quickly in our conversations and in our thought, and as a result, we find ourselves—at least in those contexts where they go by—committed to them, in just the same way that we find ourselves committed to the assertions that are made in those conversations. Given that sometimes we *learn* things in our conversations with others. Some of those things we hear and find ourselves agreeing with in conversation become standing beliefs as we come to hold onto those commitments beyond the scope of that interaction. This occurs not only through a process of considered rational reflection on each individual item that we have accepted. Sometimes we simply find ourselves continuing to believe what we were told. It is not surprising that we do so, given the function of conversation and communication, and the way we model our behaviour on others. It is not at all surprising that in a community of others who make inferences of certain kinds, we find ourselves inferring in that way, too, in what we think and say. As a result, we should expect to find ourselves committed to characterising generics, whether we have rationally reflect on those commitments or not. And upon reflection, we find that this is indeed how we behave, as thinking creatures who orient our selves to the world with views filled with generic connections and the default inferences they express.

6. OPTIONS FOR CRITIQUE & REFORM

Not all of the characterising generics in our conceptual schemes are *helpful*, or *just*, and neither are all the inferences they make explicit. Our views of others are filled with stereotypes, many of them bearing the marks of our histories as oppressor and oppressed. Our languages and conceptual schemes witness to our conflicted histories, and the language of the dominant group and the way we wield that language and our explanations of the world around us and our place in it cannot help reflect that terrain of privilege and exclusion. This goes right down to the level of the individual inferences we accept, as these are the fulcra leading from one concept to another, the tiniest links in our web of connections that represent our expectations and characterisations of what is normal and what is out of the ordinary, of what stands in need of justification, and what can count as a sufficient answer

to a given question. The links in this network of interconnections are represented explicitly by the characterising generics we accept. It is not surprising that these claims are become a site of conflict, when we disagree about whether *men are violent*, *women are submissive* or *Muslims are terrorists*. There are contested understandings of how the world *is* and how it *should be*, and it is no surprise that when we move far beyond the sharply defined quantifiable expressions that can, on occasion, be decisively settled one way or another (as much as anything is settled in the context of heated debate), to the defeasible wilds of default generalisations that are notoriously resistant to counterexample, this disagreement is well-nigh intractable.

So, what hope is there for critique and reform for contexts in which toxic and unhelpful views, expressed in generic terms, have taken root? What tools do we have at our disposal that might help elicit change? In this short concluding section, I will make some suggestions which become salient when we recognise the connection between characterising generics and default inference.

DENYING A GENERIC IS NOT ENOUGH: As we have already seen, arguing over a characterising generic *Fs are G* is true or not is unlikely to shift views concerning them. If the generic is deeply held, then no number of counterinstances, even up to granting that *most Fs aren't G* would be enough to shift commitment to *Fs are G*, as the Ross River Fever cases show us. It would be fruitful to look for an alternative strategy for shifting commitment to generic expressions.

IN LIMITED DOMAINS IT MAY HELP TO AVOID GENERICS: In a controlled environment, we can undercut the generic by being more *specific*, by moving to explicitly stated quantifiers. Instead of asking ourselves whether *Fs are G*, we can ask more sharply defined questions: are *all Fs G*? Are *most*? We can move into the mode of explicit quantification and engage in statistical inquiry, with all the rigour that entails. In certain explanatory practices, this makes a great deal of sense. It is clearer to us (to a significant degree, at least) what counts as evidence for or against an explicitly quantified claim, and if we have the syntactic discipline to restrict our vocabulary in this way, we can change the arena of dispute and engagement to this new ground. However, even if we have the discipline to restrict our vocabulary, this will not make our original conflict disappear.

THE PRACTICE OF DEFAULT INFERENCE WILL NOT GO AWAY: Since we will still engage in reasoning, in our thought and our talk, and since this is *not* (and can never be) restricted to the canons of deductive validity, we are still committed to default inference. We will still make the kinds of defeasible inference steps that ground characterising generics. If I make the default inference from *Fa* to *Ga*, then even if I refrain from articulating this with the generic expression *Fs are G*, the underlying issue (that I conceive of Muslims as terrorists, or women as submissive, or men as violent, etc.) will remain, even if I have somehow found a way to keep this out of my practice of assertion by being careful to avoid generic expressions. I still licence those inferences, and those connections are salient to me, despite having washed my thought and talk of generics. The inferences I grant remain unperturbed, and

no quietism concerning generics will shift them.

WE SEEK CONNECTIONS FOR SALIENT CONCEPTS: One way, of course, to undercut the default inference from Fa to Ga is to shift to a practice in which the question of whether something is G or not *doesn't arise*. If I do not care about whether something is G or not, and do not need to have any understanding of how G is connected to other concepts, this would allow for the inference from Fa to Ga to dissipate. If you are in a community in which Ross River Fever is not a problem, you do not form the view that mosquitos carry it. It is in those communities where people *care* about a concept that the generics become deeply rooted because we wish to understand how to avoid those things we wish to avoid, and attract those things we desire, and in any range of ways, influence our environment in ways that further our goals. In the case of unjust stereotype generics such as *Muslims are terrorists*, it is clear that this becomes salient in communities in which terrorism is a live concern, whether reasonable or not. As a result, it is unsurprising that explanatory and inferential connections take root for concepts such as these.

THE PRACTICE OF DEFAULT INFERENCE WILL NOT GO AWAY: When we *care* about the property G , we won't want to revise the concept away. So another approach is to find an alternate explanation for Ga , other than Fa . To start with our running example, if not all species or genus of mosquito carry RRF, then we could *refine* our generic to the more specific one, that *Mosquitos of genus Culex carry RRF*, while being able to also say that *mosquitos outside that genus don't carry RRF*. For communities that are able to keep track of mosquitos by specific genus, a more discriminating generic like this can easily take root and replace the less discriminating one that characterised all mosquitos 'unfairly'. Unfortunately, for a community that has no other use for distinguishing mosquitos by genus, and no ability to keep track of which kind of mosquitos are where, such a conceptual revision is unlikely to take root, as the simpler generic will be easier to work with.

To attempt this strategy to undercut *Muslims are terrorists*, we might attempt replacing them with more useful explanations which combine a higher degree of explanatory power and which use concepts that are also in wide use in the target communities. One possibility is to challenge the stereotype views of terrorists by appealing to research showing that having a history of perpetrating domestic violence is a much stronger predictor than anything else of involvement in terrorism and mass killing [31]. The concept of domestic violence is in widespread community use, and it is clearly also naturally connected with involvement in mass killing. Given that a community concerned about mass violence is looking for ways to understand it, alternative explanations using available concepts have are necessary to shift explanations from unhelpful, unjust characterisations in more just and productive directions.

NEW CONNECTIONS CAN TAKE ROOT IN ALTERNATIVE COMMUNITIES This is kind of shift takes time, and in many cases, the salient concepts are not in wide use. It took time to introduce the concept of domestic violence into wide circulation.

Given that conceptual and communicative practices need time to take root, it is unsurprising that concepts arise out of the practice of smaller communities, with their own distinctive ways of looking at the world, of drawing connections, and explaining things to themselves and to each other. Given that one way that we *acquire* our explanatory and inferential capacities is in dialogue with each other, if new connections are to emerge, they will take root in smaller subcommunities which develop not only their own new concepts, but also new kinds of salience structures, in which the inferences that are granted have shifted, and new explanations may take root.

To sum up: the distinctive behaviour of generic judgements can be explained by their grounding in our inferential and explanatory practice. Attention to that practice can aid us in our understanding the difficulties in resolving disagreements couched in generic terms, as well as pointing us to some of possibilities for reforming and revising those practices and the views they represent.

REFERENCES

- [1] N. P. ADAMS. “Authority, Illocutionary Accommodation, and Social Accommodation”. *Australasian Journal of Philosophy*, 98(3):560–573, 2020.
- [2] DAVID BEAVER AND HENK ZEEVAT. “Accommodation”. In GILLIAN RAMCHAND AND CHARLES REISS, editors, *The Oxford Handbook of Linguistic Interfaces*. Oxford University Press, 2007.
- [3] DAVID I. BEAVER AND BART GEURTS. “Presupposition”. In EDWARD N. ZALTA, editor, *The Stanford Encyclopedia of Philosophy*. CSLI, 2012.
- [4] DAVID I. BEAVER, CRAIG ROBERTS, MANDY SIMONS, AND JUDITH TONHAUSER. “Questions Under Discussion: Where Information Structure Meets Projective Content”. *Annual Review of Linguistics*, 3(1):265–284, 2017.
- [5] PAUL BOGHOSSIAN. “What is inference?”. *Philosophical Studies*, 169(1):1–18, 2014.
- [6] ROBERT B. BRANDOM. *Making It Explicit*. Harvard University Press, 1994.
- [7] ROBERT B. BRANDOM. *Articulating Reasons: an introduction to inferentialism*. Harvard University Press, 2000.
- [8] LEWIS CARROLL. “What the Tortoise Said to Achilles”. *Mind*, 4(14):278–280, 1895.
- [9] CATARINA DUTILH NOVAES. “A Dialogical, Multi-Agent Account of the Normativity of Logic”. *Dialectica*, 69(4):587–609, 2015.
- [10] CATARINA DUTILH NOVAES. “Reductio ad absurdum from a dialogical perspective”. *Philosophical Studies*, 173(10):2605–2628, 2016.
- [11] CATARINA DUTILH NOVAES. *The Dialogical Roots of Deduction: Historical, Cognitive, and Philosophical Perspectives on Reasoning*. Cambridge University Press, 2020.
- [12] KAI VON FINTEL. “What Is Presupposition Accommodation, Again?”. *Philosophical Perspectives*, 22(1):137–170, 2008.
- [13] MANUEL GARCÍA-CARPINTERO. “Accommodating Presuppositions”. *Topoi*, 35:37–44, 2016.

- [14] BART GEURTS. “Convention and common ground”. *Mind & Language*, 33(2):115–129, apr 2018.
- [15] RODERIC A. GIRLE. “Proof and Dialogue in Aristotle”. *Argumentation*, 30:289–316, 2016.
- [16] C. L. HAMBLIN. *Fallacies*. Methuen & Co Ltd, London, 1970.
- [17] SALLY HASLANGER. ““But Mom, Crop-Tops are Cute!” Social Knowledge, Social Structure and Ideology Critique”. *Philosophical Issues*, 17(1):70–91, 2007.
- [18] SALLY HASLANGER. “Ideology, Generics, and Common Ground”. In CHARLOTTE WITT, editor, *Feminist Metaphysics: Explorations in the Ontology of Sex, Gender and the Self*, pages 179–207. Springer, Dordrecht, 2011.
- [19] REBECCA KUKLA AND MARK LANCE. *Yo! and Lo!: The Pragmatic Topography of the Space of Reasons*. Harvard University Press, 2009.
- [20] MARK LANCE AND W. HEATH WHITE. “Stereoscopic Vision: Persons, Freedom, and Two Spaces of Material Inference”. *Philosophers’ Imprint*, 7(4):1–21, 2007.
- [21] MARK NORRIS LANCE AND JOHN O’LEARY-HAWTHORNE. *The Grammar of Meaning*. Cambridge University Press, 1997.
- [22] RAE LANGTON. “Accommodating Injustice—the 2015 Locke Lectures”. Recordings and handouts at <https://www.philosophy.ox.ac.uk/john-locke-lectures>, 2015.
- [23] SARAH-JANE LESLIE. “Generics: Cognition and Acquisition”. *The Philosophical Review*, 117(1):1–47, 2008.
- [24] SARAH-JANE LESLIE. “Generics Articulate Default Generalizations”. *Recherches linguistiques de Vincennes*, 41:25–44, 2012.
- [25] SARAH-JANE LESLIE. “The Original Sin of Cognition: Fear, Prejudice, and Generalization”. *Journal of Philosophy*, 114(8):1–29, 2017.
- [26] DAVID LIEBESMAN AND RACHEL KATHARINE STERKEN. “Generics and the metaphysics of kinds”. *Philosophy Compass*, 16(7):e12754, 2021.
- [27] JOHN MACFARLANE. “Pragmatism and Inferentialism”. In BERNHARD WEISS AND JEREMY WANDERER, editors, *Reading Brandom: On Making it Explicit*, pages 81–95. Routledge, 2010.
- [28] SARAH E. MURRAY AND WILLIAM B. STARR. “Force and Conversational States”. In DANIEL FOGAL, DANIEL HARRIS, AND MATTHEW MOSS, editors, *New Work on Speech Acts*, pages 202–236. Oxford University Press, New York, 2018.
- [29] BERNHARD NICKEL. *Between Logic and the World*. Oxford University Press, 2016.
- [30] BERNHARD NICKEL. “Generics”. In BOB HALE, CRISPIN WRIGHT, AND ALEXANDER MILLER, editors, *A Companion to the Philosophy of Language*, pages 437–462. John Wiley & Sons, Ltd, Second edition, 2017.
- [31] RACHEL PAIN. “Everyday terrorism: Connecting domestic violence and global terrorism”. *Progress in Human Geography*, 38(4):531–550, 2014.
- [32] F. PELLETIER AND N. ASHER. “Generics and Defaults”. In J. VAN BENTHEM AND A. TER MEULEN, editors, *Handbook of Logic and Language*, pages 1125–1179. MIT Press, 1997.

- [33] PAUL PORTNER. “The semantics of imperatives within a theory of clause types”. In K. WATANABE AND R. B. YOUNG, editors, *Proceedings of Semantics and Linguistic Theory 14*, pages 235–252. CLC Publications, Ithaca, NY, 2004.
- [34] PAUL PORTNER. *Modality*. Oxford University Press, Oxford; New York, 2009.
- [35] ARTHUR N. PRIOR. “Worlds, Times and Selves”. In PER HASLE, PETER ØHRSTRØM, TORBEN BRAÜNER, AND JACK COPELAND, editors, *Papers on Time and Tense*, pages 241–256. Oxford University Press, 2003. New edition.
- [36] GREG RESTALL. “Multiple Conclusions”. In PETR HÁJEK, LUIS VALDÉS-VILLANUEVA, AND DAG WESTERSTÅHL, editors, *Logic, Methodology and Philosophy of Science: Proceedings of the Twelfth International Congress*, pages 189–205. KCL Publications, 2005.
- [37] GREG RESTALL. “Proofnets for $S5$: sequents and circuits for modal logic”. In COSTAS DIMITRACOPOULOS, LUDOMIR NEWELSKI, AND DAG NORMANN, editors, *Logic Colloquium 2005*, number 28 in *Lecture Notes in Logic*, pages 151–172. Cambridge University Press, 2007.
- [38] GREG RESTALL. “Truth Values and Proof Theory”. *Studia Logica*, 92(2):241–264, 2009.
- [39] GREG RESTALL. “A Cut-Free Sequent System for Two-Dimensional Modal Logic, and why it matters”. *Annals of Pure and Applied Logic*, 163(11):1611–1623, 2012.
- [40] CRAIGE ROBERTS. “Information structure in discourse: Towards an integrated formal theory of pragmatics”. *Semantics and Pragmatics*, 5(6):1–69, 2012.
- [41] CRAIGE ROBERTS. “Accommodation in a Language Game”. In BARRY LOEWER AND JONATHAN SCHAFFER, editors, *A Companion to David Lewis*, chapter 22, pages 345–366. John Wiley & Sons, Ltd, 2015.
- [42] CRAIGE ROBERTS. “Speech Acts in Discourse Context”. In DANIEL FOGAL, DANIEL HARRIS, AND MATT MOSS, editors, *New Work on Speech Acts*, pages 317–359. Oxford University Press, 2018.
- [43] MANDY SIMONS. “Presupposition and Accommodation: Understanding the Stalnakerian Picture”. *Philosophical Studies*, 112(3):251–278, 2003.
- [44] MANDY SIMONS. “Foundational Issues in Presupposition”. *Philosophy Compass*, 1(4):357–372, 2006.
- [45] ROBERT STALNAKER. “Common Ground”. *Linguistics and Philosophy*, 25:701–721, 2002.
- [46] RACHEL KATHARINE STERKEN. “The meaning of generics”. *Philosophy Compass*, 12(8):e12431, 2017.
- [47] PRESTON STOVALL. “Characterizing generics are material inference tickets: a proof-theoretic analysis”. *Inquiry*, pages 1–37, 2019.