

Absolutely general knowledge*

Rachel Elizabeth Fraser¹ | Beau Madison Mount²

¹Exeter College, Oxford University

²New College, Oxford University

Correspondence

Rachel Elizabeth Fraser, Exeter College,
Oxford University.

Email: rachel.fraser@philosophy.ox.ac.uk

Abstract

There is extensive debate among contemporary philosophers about the possibility of absolutely unrestricted quantification; thus far, the debate has been almost entirely logically and metaphysically focused. We argue for a third axis of evaluation: the epistemological. We defend absolutism on epistemological grounds, by showing that one prominent and attractive alternative to absolutism—*schematism*—is epistemologically unacceptable. First, we spell out and motivate an epistemological desideratum for theories of generality, a desideratum which is easily satisfied by absolutists. Second, we consider five ways the schematist might satisfy this desideratum. We argue that none of the five ways is successful.

All things are ready, if our minds be so.

—Henry V, 4.3.71

There is extensive debate among contemporary philosophers about the possibility of absolutely unrestricted quantification (Cartwright, 1994; Williamson, 2003; Uzquiano & Rayo, 2006). Absolutists maintain that such quantification is feasible: according to the absolutist, in the right context I can use a sentence such as

(1) Everything is self-identical

to quantify over absolutely everything. In contrast, relativists maintain that only quantification over restricted domains is possible.

Konstanzer Online-Publikations-System (KOPS)
URL: <http://nbn-resolving.de/urn:nbn:de:bsz:352-2-187mmvxz3tpt5>

*We thank James Studd for useful early discussions and an anonymous reviewer for extremely helpful comments. The authors contributed equally to the paper and are listed in alphabetical order.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *Philosophy and Phenomenological Research* published by Wiley Periodicals LLC on behalf of Philosophy and Phenomenological Research Inc

The absolutist claims that, at least on some occasions, (1) can express a proposition we can write as $\langle \forall^U x x = x \rangle$: \forall^U here regimented a quantifier ranging over everything that there is or ever was or will be, everything in or out of space and time, every set and every non-set, everything concrete and everything abstract. The relativist denies that there is any such proposition: for her, in any given context, ‘everything is self-identical’ can only express some proposition $\langle \forall_1 x x = x \rangle$ involving a quantifier that is not absolutely unrestricted: for every such \forall_1 , there is a more capacious quantifier \forall_2 , encompassing everything over which \forall_1 ranges and something over which it does not.

A number of philosophers (Dummett, 1995; Glanzberg, 2004; Studd, 2019) have found relativism a natural response to Russell’s paradox, the Burali-Forti paradox, and other problems afflicting naïve set theory. But the view faces a number of problems. The best-known is the relativist’s apparent inability, by her own lights, to state her view (Lewis, 1991, 68): if relativism is true, then a claim such as ‘there is no domain that contains absolutely everything’ fails to capture the intuitive core of the relativist’s view. To express that core, one would need to express the proposition $\langle \text{absolutely no domain contains absolutely everything} \rangle$, but the relativist denies that this is possible.¹

Some relativists have responded to this problem by endorsing *schematism*—the position that absolutely general commitments can indeed be expressed, but not through quantification: instead, absolutely general commitments can be formulated using open-ended schemata, which admit of instances in arbitrary extensions of our language (Lavine, 1994; Lavine, 2006; Studd, 2019, 120–41).

This paper develops a novel objection to schematism. We argue that schematists face a distinctively epistemological challenge: they cannot make sense of the thought that we have knowledge of some absolutely general claims. In other words, the schematist can give no account of *absolutely general knowledge*. We first introduce schematism in more detail and then motivate the epistemological challenge. We then consider five schematist responses to the challenge. We argue that absolutely all of them fail.

1.1 | Schemata and schematism

In order to explain the schematist’s position, it is useful to start not with open-ended schemata but with ordinary schemata—what we term *vernacular schemata*—of the kind used in the formal sciences. Extremely simple instances of schemata arise, for instance, in standard presentations of propositional logic. Suppose that we fix a language \mathcal{L} with sentence letters P_0, P_1, P_2, \dots and a standard deductive system. ‘ $P_0 \vee \neg P_0$ ’ will be a theorem, and so will ‘ $P_1 \vee \neg P_1$ ’, and ‘ $P_2 \vee \neg P_2$ ’, and ‘ $(P_1 \wedge P_2) \vee \neg (P_1 \wedge P_2)$ ’, and so on. In order to state the law of the excluded middle, however, we do not engage in the futile attempt to write down this infinite collection of formulae. Instead, we simply write ‘ $\phi \vee \neg \phi$ ’ and let it stand for all these instances. The expression ‘ $\phi \vee \neg \phi$ ’, which is termed the *template* of the vernacular schema (Corcoran, 2006, 219–20), is not itself part of the object language \mathcal{L} : it is part of a metalanguage, consisting of the *schematic letter* ϕ joined with connectives taken from \mathcal{L} . To use it to talk about infinitely many \mathcal{L} -sentences, we need a rule—called the schema’s *side condition*—specifying what the permissible substituends for the schematic letters in the template are: in this case, the side condition is that any \mathcal{L} -sentence can be substituted uniformly for ϕ to produce an *instance* of the schema

¹In most cases the relativist will deny that this is possible because she denies that any such proposition exists; some relativists may, however, allow the existence of such a proposition but deny that it can ever be expressed. We here use ‘proposition’ (and the convention of enclosing sentences in wide angled brackets to indicate propositions) as a convenient place-holder for the semantic value of ordinary declarative sentences; no contentious ontology of propositions will be defended or assumed (except for the sake of argument in § 7).

' $\phi \vee \neg \phi$ '. So ' $P_5 \vee \neg P_5$ ', ' $(\neg P_3 \vee P_3) \vee \neg (P_3 \vee P_3)$ ', ... all count as instances of the schema. We say that a vernacular schema S holds just in case all instances of S in the object language (as specified by the side condition) are true in that language. So ' $\phi \vee \neg \phi$ ' holds, since, for every sentence ϕ of \mathcal{L} , the result of concatenating an occurrence of ϕ , a disjunction symbol, a negation symbol, and another occurrence of ϕ is a true sentence of \mathcal{L} .

In classical logic, schemata are used for many kinds of generalization—sentences (as in the example above), open formulae, names, and so on. Here, however, we shall focus on templates whose schematic letters have singular terms as their substituends: these are the only cases that are directly relevant to the schematist's attempt to achieve absolute generality. We use boldface for such schematic letters. So, for instance, where English is the base language, the vernacular schema with the template ' $\mathbf{a} = \mathbf{a}$ ' and the obvious side condition has as instances 'Simone de Beauvoir = Simone de Beauvoir', 'Greenland = Greenland', and the like.

Vernacular schemata with schematic letters for singular terms are a convenient way of expressing some general claims, but they cannot substitute for absolutely unrestricted quantificational generality: since any language we can use has only countably many singular terms, any vernacular schema can have only countably many substitution instances; thus vernacular schemata cannot simulate quantification over absolutely everything (assuming there are more than countably many things).²

To get around this problem, relativists such as Shaughan Lavine (1994, 230–32; 2006, 117–23) have invoked not vernacular but *open-ended* schemata. An open-ended schema for a language \mathcal{L} is not presented by a template in a standard metalanguage for \mathcal{L} . Instead, it is presented by a sentence in a language \mathcal{L}^+ that supplements \mathcal{L} with *full schematic variables*, which are to be interpreted as having substituends not only in \mathcal{L} (as is the case for the schematic letters in vernacular schemata) but in *any possible extension whatsoever* of \mathcal{L} .

Thus, if \mathcal{L} is English, and 'Oblagaf' is not a name in English, the expression 'Oblagaf = Oblagaf' is not an instance of the vernacular schema ' $\mathbf{a} = \mathbf{a}$ ', since the side condition permits only English singular terms as substituends for the schematic letters. But there is a possible extension of English—let us call it \mathcal{L}_1 —in which 'Oblagaf' denotes a small mountain on an as-yet-unobserved asteroid near Betelgeuse. Thus—using lower-case Greek letters for full schematic variables—the open-ended schema ' $\alpha = \alpha$ ' in \mathcal{L}^+ does have 'Oblagaf = Oblagaf' as an instance, for it allows substituends from \mathcal{L}_1 (and every other possible extension of the base language).³ An open-ended schema holds if and only if every instance, in *any* of these languages, is true: thus commitment to ' $\alpha = \alpha$ ' carries with it commitment to the self-identity of a small mountain on an asteroid near Betelgeuse; it likewise carries with it commitment to the self-identity of any object whatsoever, since, for any object whatsoever, there can be an extension of English with a new name for that object.

In stating the distinction between vernacular and open-ended schemata, we have used an absolutist idiom: a vernacular schema holds just in case every instance in the base language is true, but an open-ended schema holds just in case every instance in *absolutely* any extension of the base language is true. But this makes sense only if quantification over 'absolutely' all language extensions is possible—and the relativist denies this. Nonetheless, relativists like Lavine and James Studd (2019, 127–28) maintain that it is permissible to take open-ended schemata as primitive, independent of any notion of quantification, and that they provide a non-quantificational means to attain absolute

²The relativist could, of course, maintain that there are only countably many objects—but this commits her to a deeply revisionary metaphysics, and, plausibly, one which is in tension with our best physics.

³More formally, we can think of the meaning of an open-ended schema as given by a map from languages extending the base language to vernacular schemata in those languages (Feferman, 1991, 8).

generality. Indeed, Lavine (2006, 117) describes them as “another form of generality more primitive than quantificational generality”.

This schematist version of relativism, if it can be coherently maintained, has a number of appealing features. It would, for instance, provide a way to account for the sense that, in saying ‘everything is self-identical’, I am not leaving open the possibility of discovering some non-self-identical thing in an expansion of my current domain. Furthermore, Studd (2019, 130–32) has argued that open-ended schemata provide a stable way to state the relativist thesis (although this claim is controversial, and the details are not directly relevant to our argument).

But schematism faces difficulties. It is not clear that the fully non-quantificational account of primitive open-ended schemata required by the schematist can be made out. Furthermore, schemata are subject to severe expressive limitations: as a schema cannot be negated, there is no way for the schematist to simulate an existential, rather than a universal, absolutely general quantificational claim (Williamson, 2003, 438–39).

We put such concerns to one side. Our argument is that, however well the schematist fares in logical and metaphysical terms, she faces intractable epistemological problems.

1 | THE EPISTEMOLOGICAL OBJECTION

Consider the following claims:

- (2) Everything is self-identical.
- (3) Every set is an element of some set.
- (4) Nothing is redder than itself.

The absolutist and the schematist agree that when we make such claims, we express *absolutely general claims*, although they disagree whether the generality involves unrestricted quantification. (The schematist may hold that such claims have a semantic content that is best represented schematically, or that we successfully express commitment to some relevant schema via some pragmatic mechanism. The distinction makes no difference for our purposes.)

But in addition to this, we take ourselves to *know* the claims expressed by (2)–(4).⁴ We think you probably know them, too. And we think you can express this knowledge by uttering the relevant sentences.

This motivates:

(K-DESIDERATUM) Under the right circumstances, an agent can manifest knowledge by making an absolutely general claim.

One small clarificatory point: I can manifest knowledge that p by asserting ‘ p ’ under the right circumstances; manifesting knowledge that p does not require me to assert ‘I know that p ’—that would count as manifesting knowledge of the higher-order claim that I know that p .

The absolutist can easily meet the K-desideratum. According to the absolutist, each of (1)–(3) expresses a quantified proposition. Absolutists can simply plug in their favourite account of propositional knowledge (safety, evidentialism, virtue theory, ...) and they have a working theory of what

⁴We assume the standard iterative picture of set theory. Proponents of alternative approaches (Quine, 1937) are invited to choose another universal set-theoretical claim.

it is to know an absolutely general claim. In short, they have an easy story about what it is to have absolutely general knowledge.

Things are less rosy for the schematist. According to her, when we assert (2)–(4), we express commitment to an open-ended schema. Thus, to meet the K-desideratum, the schematist must vindicate the thought that an agent can *manifest knowledge* by expressing commitment to an open-ended schema. We argue that the schematist cannot satisfy the K-desideratum.⁵

We have thus far motivated the K-desideratum by noting that it is overwhelmingly natural to think of ourselves as knowing (2)–(4), and thus as knowing some absolutely general claims. We take the overwhelming naturalness of such a description to be excellent grounds for endorsing the K-desideratum.

But perhaps you don't think it obvious that we can manifest knowledge by making absolutely general claims. You might be wary for a few different reasons. Perhaps you're a sceptic, either across the board, or with respect to metaphysical claims in particular. In our view, such sceptics should be perfectly happy with the K-desideratum. Even sceptics about metaphysics should be happy to grant that we can know a set-theoretic claim like (3); most global sceptics should be happy to grant that we can know a claim as minimal and straightforward as (4).

But there's more to be said in favour of the K-desideratum. Contrast the following cases:

(LUCKY LOGIC LARRY) Larry has just started studying set theory. He isn't a very good student and tends to work through things too quickly. He often makes mistakes. His latest problem sheet required him to consider whether every set is an element of some set. Larry thinks about the issue briefly and decides that the answer is 'Yes'. He doesn't really have any well-worked out story about why he thinks the answer is 'Yes'—just a vague feeling that 'Yes' is the right answer. One of Larry's fellow students, Karen, asks Larry whether every set is an element of some set. Larry responds by asserting: 'Look, Karen, absolutely every set is an element of some set.'

(CAREFUL CATY) Another of Larry's classmates, Caty, decides that she should work through her problem sheet on her own. She thinks very carefully at great length about the question of whether every set is an element of some set. She realises that this follows from the standard axioms, and she considers the arguments in favour of those axioms and the problems with alternative axiomatizations of set theory on which the claim would fail. She rationally accepts the standard axioms and concludes that every set is an element of some set. Caty's fellow student Karen asks her whether every set is an element of some set. Caty responds by asserting: 'Good question, Karen. Absolutely every set is an element of some set.'

There's a clear asymmetry: Larry and Caty make the same claim. But where we judge that Caty is entitled to her assertion, we judge that Larry is not.

The absolutist can easily explain the asymmetry. According to the absolutist, both Larry and Caty assert a quantified proposition. Assertions, it is generally agreed, are norm-governed speech acts; thus, the absolutist can simply appeal to their favourite norm(s) on assertion (the knowledge norm (Williamson, 2000), the reasonable belief norm (Lackey, 2007), ...), and they have an off-the-peg account of the contrast between Larry and Caty.

For the relativist to tell a comparable story, expressing commitment to an open-ended schema must, like asserting a proposition, be governed by some norm that Caty satisfies and Larry does not. If the schematist is to explain the contrast between Larry and Caty adequately, Caty must stand in some

⁵To be clear: the schematist can tell an easy propositional story about what is involved in knowing (2)–(4) so long as they are happy to grant that none of (2)–(4) express absolutely general claims. But if they grant this, they give up on meeting basic challenges to schematism.

relation R to the schema she expresses that Larry does not stand in to the schema he expresses. Let's call this relation R *schema-knowledge*. From this set-up, the requirement that the schematist accept the K -desideratum easily falls out. The ideology of schema-knowledge also affords a nice way of stating this paper's primary thesis: no satisfactory account of schema-knowledge can be given. Thus, the schematist's account is an epistemological failure.

2 | PRIMITIVISM

THE first response that the schematist can offer to the epistemic challenge is the simplest, but also the least satisfying: the schematist can claim that schema-knowledge is a primitive mental state, incapable of being cashed out in any other terms. To be sure, sometimes claiming that a notion is primitive, capable of illustration but not analysis, is a perfectly legitimate manoeuvre.

Here, however, determined primitivism seems sorely inadequate. Compare Timothy Williamson's (2000) argument that (propositional) knowledge is primitive: we have an extremely clear pretheoretical grasp of propositional knowledge, and we can give an extensive description of its connections to other related notions (belief, safety, epistemic probability, and so on) even if we cannot give an analysis. Nothing of the sort holds for the schematist: schema-knowledge is a theoretical posit, an explication for which only becomes necessary once we admit the fairly esoteric expressive resource of open-ended schemata in order to solve the problems posed by generality relativism.

Whilst schema-knowledge is initially picked out in terms of its functional role, which parallels (in our view) that of propositional knowledge, if primitivism is true, schema-knowledge and propositional knowledge are not distinct determinates of the same determinable, or anything similar: how they come to play the same functional role is wholly unexplained. The primitivist offers us no illuminating principles connecting schema-knowledge to other properties. Presumably schema-knowledge is factive—it entails truth of all instances of the schema—but it is unclear what more could be said.

This is not a knock-down refutation of primitivism, of course, but it is clear that the burden lies on the schematist to provide some kind of motivating story for a primitivist account of schema-knowledge, and it is very hard to see what such a story would consist in.

3 | PROPOSITIONAL FUNCTIONS

The second option does provide an analysis of schema-knowledge and comes with an interesting historical pedigree. We standardly take the object of ordinary knowledge attributions to be a proposition: to know that it is raining in Paris is to know a proposition. But we might think that, alongside propositions, the world (or the type hierarchy) contains propositional functions in Russell's sense: objects which stand to open formulae as propositions stand to closed sentences. The early Russell held that these propositional functions can themselves be objects of attitudes, and linguistic expressions corresponding to them can be asserted:

When we assert something containing a real variable, as in *e.g.*

“ $\vdash \cdot x = x$,”

we are asserting *any* value of a propositional function. When we assert something containing an apparent variable, as in

“ $\vdash \cdot (x) \cdot x = x$,”

[...] we are asserting [...] *all* values [...] of the propositional function in question. [...] When we assert something containing a real variable, we cannot strictly be said to be asserting a *proposition*, for we only obtain a definite proposition by assigning a value to the variable, and then our assertion only applies to one definite case, so that it has not at all the same force as before. When what we assert contains a real variable, we are asserting a wholly undetermined one of all the propositions that result from giving various values to the variable. It will be convenient to speak of such assertions as *asserting a propositional function*. (Whitehead & Russell 1927, 1:18)

Disentangling use and mention in Russell's presentation is difficult, but the text provides the resources for an interesting variant of schematism: An open-ended schema, the Russellian schematist can maintain, is really just a formula with a free variable. Normally we do not think such a formula can be uttered assertorically on its own, but the Russellian will claim that this is just a prejudice inherited from the Tarskian tradition of equating satisfaction on all assignments with truth only for closed formulas. But, given the tight connection between knowledge and assertion, if the class of permissible kinds of asserted utterances is broader than we normally assume, it is wholly plausible that the class of objects of propositional attitudes is as well. Thus, we have a simple and easy explanation of schema-knowledge: to have schema-knowledge is to know a propositional function.

It might be objected that knowledge of propositional functions is itself an unexplained primitive notion, and the Russellian is no better off than the primitivist about schema-knowledge. But there is a substantial disanalogy: propositional functions can be naturally accommodated within the familiar hierarchy of types: the propositional function picked out by ' $x=x$ ' (or, on the schematist's view, ' $\alpha=\alpha$ '), is just a type- $(e \rightarrow t)$ entity in an intensional type theory. And replacing a single, type-specific operator (here, the type- $(t \rightarrow t)$ knowledge operator) with a family of operators taking entities of various types is one of the simplest and most natural kinds of generalization possible in type theory. For this reason, the Russellian has explanatory resources that the primitivist lacks.

But, however reasonable an account of assertion or knowledge adapted to propositional functions as well as propositions is on its own, it cannot be made to serve the schematist's generality-relativist programme. In order to gain the benefits of open-ended schemata, the schematist must maintain that absolutely unrestricted *quantification* is impossible, whereas absolutely unrestricted *schematic generality* is not. In conjunction with Russellianism, however, this yields the result that there are expressions in our language which pick out propositional functions having the entire universe as their domain, but there are no such expressions picking out quantifiers that range over the entire universe. And this is profoundly unmotivated: if complete domain-generality is there for the taking, why would the mere addition of a quantifier place it off limits? After all, if we combine the Russellian account of propositional functions with the modern understanding of quantifiers, what a quantifier picks out simply is a higher-type propositional function: a type- $((e \rightarrow t) \rightarrow t)$ function rather than a type- $(e \rightarrow t)$ function. But how can it happen that there are expressions for absolutely unrestricted type- $(e \rightarrow t)$ functions but not for absolutely unrestricted type- $((e \rightarrow t) \rightarrow t)$ functions?

If open-ended schemata are to be combined with generality relativism, then they cannot be understood as propositional functions, or as functions of any sort—for they would have to be functions with an absolutely unrestricted domain, and there is no stable position that allows us linguistic access to such functions but not to absolutely unrestricted quantifiers. Hence, the (otherwise attractive) Russell-inspired account of schema-knowledge is not available to the schematist.

4 | REDUCTIVE APPROACHES

When all goes well, we stand in the knowledge relation to propositions. Simone stands in the knowledge relation to the proposition <Greenland is cold>, and to the proposition <2 is an even number>. Call this relation *propositional knowledge*. We can also stand in knowledge relations to people and objects, and things whose ontological status—calculus, physics—whose ontological status is unclear. Simone knows Jean-Paul, she knows Paris, and she knows calculus. As Sophie-Grace Chappell notes (2012, 185): “When I have objectual knowledge, say of the tree in my garden, what I know is *the object*, the tree: not some proposition about the tree, or some experience of the tree, or some technique relevant to the tree”. Call this relation *objectual knowledge*.

One simple way for the schematist to satisfy the K-desideratum would be for her to show that having schema-knowledge can be reduced to having propositional knowledge or to having objectual knowledge:

(PROPOSITIONALISM) To have schema-knowledge is to have propositional knowledge.

(OBJECTUALISM) To have schema-knowledge is to have objectual knowledge.

Later in the paper, we will make some minimal, and (we think) highly plausible assumptions about what propositional knowledge requires; for now, we make no such assumptions: our rejection of propositionalism should be acceptable regardless of your epistemological proclivities.

4.1 | Propositionalism

Let \mathcal{S} be a name for the open-ended schema, commitment to which we express or aim to express when we assert (1). There are two different propositionalist strategies. Option one says: having schema-knowledge of \mathcal{S} is a matter of knowing the propositions which are \mathcal{S} 's instances. Option two says: having schema-knowledge of \mathcal{S} is a matter of knowing something *about* \mathcal{S} .

Option one faces an obvious difficulty. If schema-knowledge of \mathcal{S} requires knowing all the propositions that are expressed by instances of \mathcal{S} , then, because \mathcal{S} is open-ended, the relativist risks being able to give an account of schema-knowledge only by covertly helping themselves to absolutist resources. (We consider a sophisticated rejoinder in § 5.) A weakened version of this first option, on which schema-knowledge requires knowledge only of *some* of the propositions expressed by a schema's instances, looks unappealing for different reasons.

Knowing some of the propositions expressed by a schema's instances cannot be sufficient for schema-knowledge. If it were, I would count as schema-knowing the open-ended schema ‘ α is is something I have touched’. So construed, either schema-knowledge cannot play its key functional role, or I am entitled to assert ‘I have touched absolutely everything’.⁶

Option two faces similar difficulties. First, it involves *intellectualizing* absolutely general commitments. Our first-year undergraduates know that everything is self-identical. But they don't know anything about schemata. So it is very odd to attribute to them the knowledge of some proposition about \mathcal{S} . Second, it is very difficult to specify any proposition that might play the role the schematist requires. The obvious candidates—that all or some of \mathcal{S} 's instances are true—are non-starters for the same reasons that doom option one.

⁶We consider a more sophisticated relative of this view in our final section.

Alternatively, the relativist might introduce some new piece of vocabulary. They might, for example, introduce the notion of a *healthy*, as opposed to an *unhealthy* schema, and say that what an agent with schema-knowledge knows is that the schema is healthy. But in the absence of some interpretation of this new vocabulary, its introduction is incapable of doing any explanatory work.

4.2 | Objectualism

Can the objectualist do better? Objectualism is slipperier than propositionalism, because objectual knowledge is less well-theorised than its propositional counterpart. On one tempting analysis of objectual knowledge, it just is a species of propositional knowledge. Given the discussion above, no such reductive story will help the schematist.

There are two alternatives. The first reads objectual knowledge as acquaintance: to know Jean-Paul is to be *acquainted* with him. The second reads objectual knowledge as something like competence or know-how: to know Paris is to know how to find one's way around Paris. For now, we put the second suggestion to one side; we deal with it in § 6. That leaves us with the acquaintance reading.

John Hawthorne and David Manley (2012) consider three ways of making the notion of acquaintance precise:

(EXISTENCE) To be acquainted with an object *o*, one must know that *o* exists.

(CAUSAL) To be acquainted with an object *o* one must be causally related to *o* in some way.

(DISCRIMINABILITY) To be acquainted with an object *o*, one must know which object *o* is.

EXISTENCE both over- and undergenerates schema-knowledge.

First, suppose that for any schema *S* such that I know *S* exists, I count as acquainted with *S* and thus as having schema-knowledge of *S*. Then I know that the following schema exists: '*α* is red'. But surely I don't know that everything is red. EXISTENCE also undergenerates schema-knowledge. Consider Nominalist Nellie, who doesn't believe linguistic objects exist. Assuming that knowledge that *p* requires belief that *p*, Nominalist Nellie does not know, of any schema, that it exists.⁷ But she can surely nonetheless know that absolutely everything is at least as red as itself.

CAUSAL may be thought a non-starter: although we can stand in causal relations to *inscriptions* of schemata, we arguably cannot stand in causal relations to schemata (syntactic types, not tokens) themselves. If we can stand in causal relations to schemata, however, CAUSAL overgenerates for the same reasons as EXISTENCE: if I am in causal contact with '*α* = *α*', I am surely also in causal contact with '*α* is red'.

As for DISCRIMINABILITY, in order to assess its prospects, we need to make sense of the ideology of 'knowing-which' or 'knowing-what' (see Dummett 1991, 126–31) to which it appeals. We focus on a proposal due to Gareth Evans (1982, 107–12), which explains knowing-which in terms of *fundamental ideas*. A fundamental idea of an object *o* is a concept which encodes a property possessed uniquely by *o* and which explains what makes *o* different from everything else. We know which object *o* is, Evans thinks, so long as we have some idea (not itself necessarily fundamental) *i* of *o* such that

⁷The assumption that knowledge requires belief is widely accepted; for a defence of this orthodoxy, see Rose & Schaffer (2013). But even the heterodox can agree with our verdict on Nellie, so long as they are happy to allow that an agent knows that *p* only if they do not believe the negation of *p*.

we know what it would take for any thought built out of i , any arbitrary fundamental idea, and an idea of the identity relation, to be true.

Now, given certain ways of thinking about fundamental ideas—for example, accounts on which fundamental ideas encode properties of spatio-temporal location—it seems impossible to be acquainted in the relevant sense with an object like a schema. We don't think we know what it would take for the thought $\langle S$ is identical with the leftmost object on my desk right now \rangle to be true. It's a pretty weird thought, after all.

Let's assume, though, that we can make sense of the relevant thoughts. I will count as being acquainted with a schema S so long as, given any fundamental idea i , I have some way of thinking about S , i_S , such that I know what it would take for the thought that i is identical with i_S to be true. It seems perfectly clear that I might satisfy such a condition with respect to S without being licensed to express commitment to it. Let's make this vivid. If I count as acquainted with *any* schema, I count as acquainted with the schema ' α is red'. But I'm not licensed to express commitment to that schema! Thus the Discriminability option overgenerates schema-knowledge.

5 | THE DOUBLING-DOWN OPTION

Reductive and Russellian approaches both attempt to explain schematic knowledge by appealing to non-schematic ideology. But what if the schematist simply attempts to explain schema-knowledge using open-ended schemata? We term this the *doubling-down option*. In contrast to the propositional version of the reductive approach, the doubling-down approach takes schema-knowledge to be explicable as knowledge of each member of an open-ended cluster of propositions over which we can generalize only schematically.

Structurally, there is nothing illegitimate about the doubling-down option: if the schematist is correct, open-ended schemata are a legitimately primitive logical device, and it would be unprincipled for the anti-schematist to admit open-ended schemata as ways of expressing non-epistemic claims but to discountenance them in epistemic cases.

We argue, however, that the doubling-down option is unavailing. We shall consider a few ways of spelling out the doubling-down option, starting with the simplest.

5.1 | Straightforward doubling-down

On this account, schema-knowledge of $\alpha = \alpha$ is explained through the schema ' $K_S \alpha = \alpha$ ' (where K_S is the knowledge operator for the relevant agent S). According to a schematist of this kind, what it is for an agent to have schema-knowledge is for that agent's open-ended commitment to each instance \mathbf{I} of the schema to yield, for each \mathbf{I} , knowledge that \mathbf{I} .

This can be put more picturesquely. Just as most epistemologists endorse an account of propositional knowledge on which an agent knows that p if they have the *right kind of belief*—perhaps belief which is supported by evidence (Feldman & Conee, 1985), belief which is safe (Williamson, 2009; Pritchard 2009), belief which meets a no-defeat condition (Lackey, 2009, 44; Pryor, 2013), or belief which is supported normically (Smith, 2017)—that p , the schematist who adopts the straightforward doubling-down option takes schema-knowledge to be a matter of having the right kind of open-ended commitment to a schema. And they take an agent to have the right kind of commitment to a schema just when each of its instances is known. It is important to

distinguish this piecewise state from knowledge that *every* instance is true—knowledge of a single proposition about all the instances—through which some propositionalists attempt to explain schema-knowledge.

The problem with this straightforward response is that the open-ended schema ' $K_S\alpha = \alpha$ ' (holding fixed an appropriate S) appears to have false instances. Consider John, a competent logician: if anyone knows that everything is self-identical, then John does. But let Π abbreviate (in our metalanguage) a million-digit numeral in John's language. John cannot process Π , as it vastly exceeds his capacities of memory and cognition; *a fortiori*, he cannot process $\Pi = \Pi$, so he cannot believe, and thus cannot know, that $\Pi = \Pi$. Furthermore, even if John could form the relevant belief, he could not distinguish the true proposition $\langle \Pi = \Pi \rangle$ from very similar false ones: if Π' abbreviates a million-digit numeral that differs from Π only in its 536217th place, John will not be able to distinguish between $\langle \Pi = \Pi \rangle$ and $\langle \Pi = \Pi' \rangle$, and thus will not count as having a safe belief.

5.2 | Dispositionalism

The proponent of doubling-down is not out of options. They can tweak their account. The straightforward doubler-down identifies schema-knowledge with the right kind of commitment to a schema, and takes this 'right kind' of commitment to be a matter of each of the schema's instances being known. A more sophisticated version of the doubling-down option involves hanging on to the thought that schema-knowledge is a matter of having the right kind of commitment to a schema, but, rather than requiring the schema-knower to know each of its instances occurrently, requires only that each of its instances be such that the schema-knower has *dispositional* knowledge of it. Thus, instead of ' $K_S\alpha = \alpha$ ', the sophisticated doubler-down represents schema-knowledge using the schema ' $\diamond K_S\alpha = \alpha$ ', where \diamond is an (agent-relativized) disposition operator.

The straightforward doubler-down could not accommodate the thought that John knows that everything is self-identical, because ' $K_{\text{John}}\alpha = \alpha$ ' appeared to have false instances. It is much less obvious that ' $\diamond K_{\text{John}}\alpha = \alpha$ ' has false instances. It's obvious that John does not know that $\Pi = \Pi$; it's much less obvious that he is not *disposed* to know it.

Broadly speaking, there are two classes of circumstances under which an agent is disposed to know something without in fact knowing it. In the first sort of case, an agent who is disposed to know p fails to know p because the relevant activation conditions for the disposition are not present. Compare: a fragile vase is one with a disposition to break *when dropped or struck*. But a fragile glass may avoid breaking so long as the activation conditions—impact or application of force—never obtain. To apply this model to John would be to say that one of the activation conditions for knowing $\Pi = \Pi$ is John's grasping or entertaining the thought $\langle \Pi = \Pi \rangle$ and that, were these conditions to be activated, he would come to know that $\Pi = \Pi$.

In the second sort of case, an agent who is disposed to know p fails to know p because the disposition is finked or masked. Compare: a vase's disposition to break is *finked* where, if the activation conditions (say, being dropped) had kicked in, something (for instance, a benevolent magician who swiftly alters the vase's molecular structure immediately before it hits the floor) would have caused the manifestation not to eventuate. In this case, the vase remains fragile even though it would not break if dropped. If the vase would fail to break because it is surrounded by foam packaging, its disposition has been *masked*. To apply this second model to John is to say that John has a disposition to know p that has been finked or masked. We discuss these two options in the next sections.

5.3 | Activation accounts

Clearly, for the first strategy to be illuminating, we need to say something about what the activation conditions for knowing are. There are two obvious relevant candidates: reflection—merely turning one’s mind to a subject matter—and the stronger condition of *grasping* a proposition. Reflection is an appealing candidate for ordinary cases of dispositional knowledge: a diner in a café may not know how many chairs there are at the table at the edge of her field of view, for she has never pondered it, but she will acquire the knowledge as soon as she turns her mind to the subject. But reflection, in this weak sense, will not help the doubler-down: even if John is presented with a (very long) piece of paper containing the decimal expansion of Π and turns his mind to it, he will not know that $\Pi = \Pi$. So the doubler-down must say that the activation conditions for knowing involve grasping. But this both over- and under-generates dispositions to know.

Positing grasping as an activation condition for knowledge seems to make dispositional knowledge too easy to come by. Consider the situation of a woman, Pattie, in 1950s America who is routinely sexually harassed in the workplace. She lacks the concept *sexual harassment*; accordingly, she is not able to grasp the proposition $\langle I \text{ am being sexually harassed right now} \rangle$ (Fricker, 2007, 149–52). However, were she to acquire the concept, and so come to grasp the proposition $\langle I \text{ am being sexually harassed right now} \rangle$, she would immediately come to know it—after all, it’s perfectly obvious to anyone who is competent with the concept that what she was enduring is sexual harassment. Now, if we allow that grasping a proposition is an activation condition for knowledge, we are under pressure to allow that Pattie is disposed to know that she is being sexually harassed. But this seems like the wrong result.

Grasping also makes dispositional knowledge too hard to come by. Consider the scientist Mary from Frank Jackson’s (1982) knowledge argument. If dualism is true, then she cannot grasp propositions about redness while she remains in her black-and-white room, for she can only acquire familiarity with redness by phenomenal acquaintance. Nonetheless, she has read about redness in her textbooks on colour vision, and this seems intuitively sufficient for her to come to know the simple truth that redness is identical to redness—even if she does not know what redness is like. Of course, dualism may well be false—but the adequacy of an account of dispositions to know, presumably applicable to all rational agents in all worlds, should not be hostage to the details of the metaphysics of human minds.

Moreover, the account in terms of grasping may not even deliver all the counterfactuals the doubler-down needs. The doubler-down needs the following counterfactual to be true: were John to grasp the claim $\langle \Pi = \Pi \rangle$ then he would know it. But plausibly, the closest world in which John grasps $\langle \Pi = \Pi \rangle$ is a world w in which his cognitive capacities have been expanded *just enough* for him to grasp $\langle \Pi = \Pi \rangle$, but no more. But we find the following claim appealing:

ABILITY CONDITION An agent knows that p only if they have a reliable ability to grasp the proposition that p .⁸

Not every world in which an agent ϕ -s is a world in which they exhibit a *reliable ability* to ϕ : someone who wins the lottery by getting lucky does not have a reliable ability to win the lottery. Now consider the following question: if the actual world is one in which John cannot grasp the proposition that $\langle \Pi = \Pi \rangle$, is the closest possible world in which he does grasp this proposition a world in which he does so by getting lucky, or one in which he does so in virtue of a reliable ability to grasp it? We think

⁸Note that this neither implies nor is implied by standard reliabilist or safety conditions.

the former a more natural answer. Consider an analogy: Tushar currently lacks the reliable ability to do a cartwheel. The closest possible world in which Tushar performs a cartwheel is one in which he pulls one off by getting lucky, not one in which he has acquired a reliable ability that he in fact lacks. So we are under some pressure to think that the closest worlds in which John grasps $\langle \text{II} = \text{II} \rangle$ are worlds in which he does so without having a reliable ability to do so, and are thus worlds in which he fails to know $\langle \text{II} = \text{II} \rangle$.

5.4 | Finks and masks

The second strategy which the doubler-down can invoke is to claim that, even though some counterfactuals about knowledge of instances of the schema under the activation conditions are false, an agent with schema-knowledge is nonetheless disposed to know each instance. It is normally accepted that the counterfactual analysis of dispositions is only an approximation: a disposition can subsist even though, if its activation conditions were to obtain, it would be finked or masked (Martin, 1994; Lewis, 1997; Fara, 2005).

The doubler-down who appeals to these resources will claim that, in every case where an agent with schema-knowledge would fail to know an instance of the schema even if the activation conditions obtained, the agent's disposition to know is not absent: it is merely finked or masked.

To do so, however, the doubler-down must invoke dispositions with far more pervasive, and far more extreme, finking and masking than the paradigm cases of dispositional abilities. Let's start by getting a feel for the sort of work that we can reasonably expect finks and masks to do for us.

Here's an illustrative example: Dispositional accounts of the mind often appeal to finks and masks to deal with the rule-following puzzle: my ability to follow the rule for addition, although I would make addition errors given sufficiently large numbers, is explained by the fact that I retain the disposition to add even in the cases in which I would fail, although my disposition is finked (Martin & Heil, 1998). In such cases, finks and masks are used to extend a well-behaved rule from finitely many cases where the counterfactual would hold to at most countably many cases where it would not.

Finks and masks would need to play an extraordinarily different role if they are to do the kind of heavy lifting required by sophisticated dispositionalists.

Recall that, in order for an open-ended schema—the kind of schema to which the schematist appeals—to hold, any instance *in any expansion of the language whatsoever* must hold. A language here is an interpreted language, so there are at least as many expansions as there are interpretations of a countable collection of new constants—and thus at least as many as there are objects. Even if a relatively hardline version of nominalism is true, contemporary physics tells us that there are more-than-continuum-many spacetime regions, and thus more-than-continuum-many objects. Among those languages, there are ones whose new terms denote only extraordinarily gerrymandered regions with which we have no causal contact. There are ones whose new terms denote only enormous natural numbers—not in the orderly way that the system of decimal numerals does, but in a completely haphazard and *ad hoc* fashion. There are even languages with nonrecursive syntax—languages where it is not possible to apply a decision procedure to determine whether a particular expression is well-formed or not. (In fact, the recursive languages will be measure zero in the space of all languages.) In such cases, no actual agent will have a chance of ever occurrently knowing the vast majority of the instances of the relevant schema.

One might plausibly invoke finks and masks if they had only to bridge the gap between finitely many cases in which an agent's disposition to know would be manifested and countably many cases with the same basic structure in which the disposition would not be manifested. But here finks and

masks are called on to bridge the gap between the finite number of manifestation scenarios and a heterogeneous panoply of cases in a more-than-continuum-sized array of languages, almost all of them lacking anything like an orderly structure that a human can recognize. If finks and masks are constrained in any way by the kinds of limitations found in the paradigm cases—and it is hard to see how we could come to possess the concept if they were not—they cannot do this work.

5.5 | The hostage problem

There are yet more problems for the sophisticated dispositionalist. These problems arise for both vernacular and open-ended schemata; for the sake of simplicity, we focus on the former when possible.

Consider the situation of Parsimonious Pete. Parsimonious Pete is strongly committed to the claim that everything is self-identical, and we would normally be inclined to describe him as knowing that everything is self-identical.

Now, Parsimonious Pete does not believe that fictional characters exist. He is also a negative free logician and thus thinks that the sentence ‘Pegasus = Pegasus’ is true only if Pegasus exists. Accordingly, he does not believe that Pegasus = Pegasus. Assuming knowledge entails belief, it follows that Pete does not know that Pegasus = Pegasus.

But now suppose that in fact, Parsimonious Pete is mistaken as to whether fictional objects exist, and that ‘Pegasus’ *does* refer to an object. If Pegasus does in fact exist, then ‘Pegasus = Pegasus’ is a genuine instance of the schema ‘**a = a**’. Thus there is an instance of the schema ‘**a = a**’ of which Pete clearly lacks even dispositional knowledge. If knowing that everything is self-identical requires dispositional knowledge of each instance of ‘**a = a**’, then Pete does not know that everything is self-identical.

This is a very peculiar result: Pete’s mistaken views about the metaphysics of fictional characters appear to prevent him from knowing the most basic of logical truths! Call this the *hostage problem*: Pete’s knowledge that everything is self-identical is hostage to his false beliefs about the ontology of fictional characters. This is bad: someone ought to be able to know that everything is self-identical whether or not they have the right metaphysics of fictional objects.

The hostage problem is perfectly general. If I mistakenly think that Homer does not exist, and that Homer = Homer only if Homer exists, then I will reject the claim that Homer = Homer. But given that my beliefs about Homer are incorrect, and ‘Homer = Homer’ is an instance of the schema ‘**a = a**’, I will, by the doubler-down’s lights, fail to know that everything is self-identical. Here, my knowledge that everything is self-identical is hostage to my false beliefs about Homer. But that is absurd: someone can have an incorrect theory of the authorship of the Greek epics, and have basic metaphysical knowledge.

Importantly, the hostage problem arises for a schematist regardless of her view as to which values we might assign to schematic letters of *S* in order to produce a genuine instance of a schema *S*. A schematist might say that, whether or not ‘Pegasus’ refers, ‘Pegasus = Pegasus’ counts as an instance of ‘**a = a**’. Call this the liberal view. Liberal views may be extreme or moderate, where the extremist says that ‘Pegasus = Pegasus’ is true regardless of whether ‘Pegasus’ refers, and the moderate says that, if ‘Pegasus’ does not refer, ‘Pegasus = Pegasus’ is either false or lacks a truth value. In contrast, on a conservative view, ‘Pegasus = Pegasus’ counts as an instance of ‘**a = a**’ only if ‘Pegasus’ refers.

For both conservative and liberal, so long as there is some name **a** which in fact refers but which Pete falsely believes does not refer, there will be some (true) instance of the schema ‘**a = a**’ which Pete disavows. Pete will thus count as failing to have schema-knowledge of this schema.

At this point, the schematist might want to strengthen her account. Perhaps schema-knowledge of *S* requires not only a disposition to know all of its instances, but competence with *S*’s side condition.

In other words, perhaps schema-knowledge requires an ability to know, for any putative instance of S , whether it is a genuine instance.

The conservative cannot in good conscience impose any such constraint: to meet it, an agent would need to know, for any putatively referring term, whether it refers or not. This would automatically bar anyone with false views about which objects exist from schema-knowledge of $\mathbf{a} = \mathbf{a}$.

Nor does such a strategy help the liberal. Parsimonious Pete can recognise that ‘Pegasus = Pegasus’ is an instance of the schema ‘ $\mathbf{a} = \mathbf{a}$ ’. He just won’t accept that the instance is true. Accordingly, a Parsimonious Pete who is, by the lights of the liberal, competent with respect to the schema’s side condition will simply fail to exhibit any open-ended commitment to it, for he will take it to have false instances. Such a version of Parsimonious Pete lacks the schema analogue of belief, as well as the schema analogue of knowledge.

The last hope for the liberal is to restrict their attention to agents who are both competent with respect to the side condition *and* accept that ‘ $\mathbf{a} = \mathbf{a}$ ’ holds regardless of whether \mathbf{a} refers. We take such a restriction to be highly unmotivated and deeply costly: if the liberal relativist can grant schema-knowledge of $\mathbf{a} = \mathbf{a}$ only to agents who reject negative free logic, her theory of schema-knowledge is severely limited. But there are even worse problems. Even if such a strategy delivers the right results for ‘ $\mathbf{a} = \mathbf{a}$ ’, it fails when confronted with other schemata.

Consider the following open-ended schema: ‘ α actually exists’. The liberal—who, recall, allows that we may, by substituting non-referring terms for schematic variables, produce genuine schema instances—is committed to such a schema having some genuine instances which fail to be true.⁹

Thus, the liberal must identify schema-knowledge of ‘ α actually exists’ not with knowledge of each instance of the schema (which would be impossible) but with knowledge of each *true* instance of the schema. But such knowledge will be very hard to come by. Consider a physicist, Tilly who is doing an experiment, which may or may not result in the production of a particle, which she has decided to call ‘Harry’. If Harry is produced, she will see a distinctive trail of ionized gas particles. Suppose that, in fact, her experiment is successful, and she forms the true belief that Harry actually exists. Unfortunately for Tilly, her equipment is malfunctioning: in those close worlds in which the experiment was unsuccessful, the cloud chamber mimics the appearance of a trail of ionized gas particles. There are thus close worlds in which Tilly believes that Harry actually exists and this belief is not true. Thus, given a minimal safety constraint on knowledge, in the actual world, Tilly does not know that Harry actually exists.¹⁰

There is thus a true instance of the schema ‘ α actually exists’ which Tilly fails to know, despite being competent with all the relevant side condition, and having all the right meta-linguistic views. Even if we spot the liberal schematist an immense number of *ad hoc* restrictions, she still cannot get her view to work.

It is worth being clear that absolutists are not afflicted by the hostage problem. On any viable view of how natural language quantification works, agents can—and often do!—successfully quantify over objects whose existence they deny. If I sincerely but mistakenly deny that there are mice in my

⁹We assume here both that at least some possible but non-actual object is nameable (which strikes us as uncontroversial) as well as (more controversially) that necessitism is false. A necessitist might insist that only possible objects are possibly nameable and that, given that all possible objects exist, all instances of the schema are true (Williamson, 2013). Necessitists should turn their attention to the schema: ‘ α is actually concrete’. Our remarks will apply *mutatis mutandis*.

¹⁰If you prefer no-defeat conditions on knowledge to safety constraints, suppose that Tilly receives misleading but authoritative testimony that her experiment failed to produce Harry, but dogmatically sticks to the belief that Harry was produced. Here, Tilly fails to know a true instance of the relevant schema because her belief goes against her evidence, rather than because of her belief’s modal profile.

house, I say something false if I insist ‘There are no mice in my house’. And no matter how staunchly Parsimonious Pete denies that fictional objects exist, he does not succeed in speaking truly if he intends to make an absolutely general claim, and asserts ‘Absolutely nothing is a fictional object’.

5.6 | The conditional strategy

A schematist might reply to the hostage problem that, in cases such as that of Parsimonious Pete, Parsimonious Pete does not have schema-knowledge of ‘ $\alpha = \alpha$ ’.

Perhaps Pete should be said only to have schema knowledge of the weaker ‘ α exists $\rightarrow \alpha = \alpha$ ’. Pete will affirm ‘if Pegasus exists, Pegasus = Pegasus’ whether or not he believes that Pegasus exists; he will affirm it even if he has no views whatever about Pegasus’s reality. If this is the case, empty names will not pose a direct undergeneration problem in cases like Pete’s for the sophisticated dispositionalist. Let’s call this the *conditional strategy*.

Proponents of the conditional strategy face two problems. Consider the case of Renata, who has dispositional knowledge of every instance of $\alpha = \alpha$. Renata, then, has schema-knowledge of $\alpha = \alpha$. She is, then, presumably entitled to assert ‘Everything is self-identical’ and expresses her schema-knowledge of $\alpha = \alpha$ when she makes such an assertion. The problem is this: it’s appealing to think that if Renata and Parsimonious Pete both utter the sentence ‘Everything is self-identical’, and both thereby express knowledge, they express knowledge of the same thing. But this cannot be so: Parsimonious Pete does not know the unconditional schema; he and Renata must be expressing different states when they utter ‘Everything is self-identical’. That’s an odd result.

Even worse, the conditional strategy also overgenerates schema-knowledge. Consider the case of Prudent Phyllis. Phyllis *explicitly* disavows the sentence ‘Everything is self-identical’. Indeed, she maintains that such a sentence expresses a falsehood. But she is convinced of a weaker thesis: that all *nameable* objects are self-identical, but that there are objects which are necessarily unnameable. Phyllis—an extreme liberal—is then disposed to accept an instance of ‘ α exists $\rightarrow \alpha = \alpha$ ’ in any extension of her language. Anyone who insists that Parsimonious Pete knows that everything is self-identical in virtue of having schema-knowledge of the conditional schema must admit that Prudent Phyllis also knows that everything is self-identical. But this is clearly false. We conclude that the conditional strategy is not successful.

6 | THE PRACTICAL APPROACH

One might be tempted by something like the following picture: schema-knowledge of some schema S is a matter of knowing how *to do something* with or to a schema.

There are two available positions with respect to know-how: intellectualism and anti-intellectualism. Intellectualists take know-how to be just another species of propositional knowledge. For the intellectualist, to say that Simone knows how to make an omlette is to say that there is some proposition of a form akin to <this is a way to make an omelette> (where ‘this’ picks out the way demonstratively) that Simone knows *under a practical guise* (Stanley & Williamson 2001; Stanley 2011). Non-intellectualists reject this, and identify know-how with a form of ability or competence which, it is said, cannot be reduced to knowledge of propositions (Ryle 1945; Noe 2005).

We are both sympathetic to intellectualism, but do not assume it in what follows—we do not want our case against schematism to hang on a particular view about know-how.

Consider, first, the simple practical view on which schema-knowledge of S is a matter of knowing how to manipulate or fill in a schema. Such a view will both under and over-generate schema-knowledge. Sally may be completely hopeless when it comes to manipulating devices like schemata, but know that everything is self-identical. And Cressida may be highly competent when it comes to manipulating a schema S of which it is deeply implausible that she have schema-knowledge. Consider: Cressida knows how to manipulate both of the following schemata:

(5) $\alpha = \alpha$

(6) $\alpha \neq \alpha$

If schema-knowledge of (5) is to license an assertion that everything is self-identical, schema-knowledge of (6) must license an assertion that everything is not self-identical. But no one can be in a state in which they are licensed to assert that everything is self-identical and also that everything is not identical with itself. Thus, the simple practical view is inadequate.

One might try to block these worries by going for a hybrid view, on which schema-knowledge requires both knowledge of how to manipulate the schema and knowledge of the schema's true instances. But such a patch is not very attractive: not only is such a hybrid approach still prone to under-generate schema knowledge in just the same way as the simple view it is designed to improve upon, the condition will be trivially satisfied by any schema—like ' $\alpha \neq \alpha$ '—which has no true instances. One could in turn try to control for this by imposing the constraint that if an agent has schema knowledge of some S then S has at least one true instance. But this won't do either. Consider the schema ' α is red'. Such a schema has true instances, and I might know, for each true instance \mathfrak{I} , that \mathfrak{I} . But I am not entitled to assert that absolutely everything is red! The requirement would need to be strengthened: we must require that all the schema's instances are true and known. This iteration of hybridism, in effect, combines a requirement that one know how to manipulate the schema with some version of the doubling-down strategies explored above. We take ourselves to have shown that doubling-down strategies are not promising; they are no more promising when supplemented with a know-how requirement.

A different practical approach might identify schema-knowledge with knowing, of each instance of the schema, how to respond to it. This does not strike us as promising either. This can be given either an intellectualist or an anti-intellectualist reading. When given the former reading, it amounts to something like the following view: for each instance of the schema, an agent with schema-knowledge knows whether to accept that instance. For now-familiar reasons, this will not work. Such knowledge must be either occurrent or dispositional; neither option avails. And someone like Parsimonious Pete, for example, will fail this test for schema knowledge. When given an anti-intellectualist reading, the view looks even more unappealing: it is highly implausible that ordinary agents have anything like this ability.

7 | DO OUR OBJECTIONS PROVE TOO MUCH?

There is, however, an important worry about these responses to the doubling-down and practical proposals that merits separate treatment. It might be thought that our arguments prove too much if they prove anything.

The defender of schematism can argue as follows: *everyone* makes use of inference rules such as modus ponens and universal instantiation. These rules play an essential role in deductive reasoning: by making use of them, we expand our stock of knowledge. But using these rules can only be a way of expanding our knowledge if we know them—and that knowledge is not simply a matter of knowing

some proposition, as Lewis Carroll's (1895) regress argument showed.¹¹ Instead, the objector will continue, knowing modus ponens is a kind of rule-knowledge, akin in this respect to the schema-knowledge that we maintain the relativist needs to be able to explain. But clearly an agent can, for instance, know a rule of inference even whilst being disposed to make mistakes about (for example) what counts as a genuine instance of the rule. So our requirements for schema-knowledge are unduly demanding: they predict that we lack knowledge of rules such as modus ponens.¹²

We find this response unavailing. It is useful to distinguish two puzzles about inference rules. The first—which we call the *form puzzle*—derives from the fact that a rule has premises and a conclusion, and it in some sense corresponds to a movement from the former to the latter, whereas a single proposition does not. We can distinguish, for instance, between the rule instance of modus ponens represented by the sequent $P_0, P_0 \rightarrow P_1 \Rightarrow P_1$ and the corresponding conditional represented by the sequent $\Rightarrow (P_0 \wedge (P_0 \rightarrow P_1) \rightarrow P_1)$; Carroll's Tortoise accepts the latter but not the former. The second puzzle—which we can call the *rule-generality puzzle*—reflects the fact that knowledge of a rule extends beyond knowledge of any particular instance of the rule.

The form problem is a difficult problem for every account of logical knowledge, and we do not propose to solve it here; we assume that whatever the right solution is will be available to generality absolutists and generality relativists on equal terms. The rule-generality problem, in contrast, is structurally similar to the problem of accounting for schema-knowledge, and reduces to it in the case of zero-premiss rules.

The objector proposes to handle the rule-generality problem by adopting *imperfect dispositionalism*. On their view, agents count as knowing a rule so long as their dispositions to know conclusions of instances of the rule, given knowledge of the premisses, are *good enough*. The fact that there are some failures of dispositional knowledge does not prevent one from knowing a rule, on pain of an unpalatable scepticism about ordinary agent's knowledge of rules like modus ponens.

But we do not think that imperfect dispositionalism, on its own, is a plausible solution to the rule-generality problem. Consider the following case:

(JAMIE'S LOGICAL REVISIONISM) Jamie has near perfect dispositions when it comes to modus ponens. She is disposed to know Q whenever she knows P and $P \rightarrow Q$. However, she has recently formed the unusual view—based on her speculations in the philosophy of physics—that modus ponens is truth-preserving only in instances where the premisses do not involve names of actually existing tachyons. As it turns out that there are no actually existing tachyons, this does not affect any instances of her reasoning. But if asked, she will maintain that modus ponens, in full generality, is not a valid rule of inference.¹³

Does Jamie know modus ponens? We think the answer is clearly 'no'. Nonetheless, on any reasonable construal, Jamie's dispositions are better than those of the ordinary reasoner: where the ordinary reasoner occasionally makes errors in the use of modus ponens, Jamie does not. So an anti-sceptical stance with respect to the ordinary reasoner's knowledge of rules like modus ponens does not motivate imperfectionist dispositionalism. This does not conflict with the fact that Jamie can extend her knowledge using modus ponens: anyone with externalist sympathies will be happy to allow that use

¹¹On this argument, see generally Besson (2018).

¹²We thank the anonymous referee for pressing us on this point.

¹³Compare the discussion of Vann McGee's (1985) in Williamson (2007, 85–98).

of a good rule under the right circumstances can be knowledge-extending even if we don't know the rule itself (Phillie, 2007).

The absolutist, on the other hand, has a natural story to tell about our knowledge of modus ponens: we know modus ponens because we know that absolutely every instance of modus ponens is valid. (We know this, perhaps tacitly; Jamie doesn't, tacitly or otherwise.) The absolutist need not claim that knowledge of modus ponens can be *reduced* to absolutely general propositional knowledge, since the form problem stands in the way; nonetheless, knowledge of a rule has a propositional *component* that he can account for easily.

The *strident absolutist* insists that only she can tell such a story; after all, she can quantify over every instance of modus ponens; her schematist opponent cannot. A more conciliatory absolutist allows that the schematist can tell an augmented version of this story. To know that every instance of modus ponens is valid does not require absolutely general quantification, because it only requires quantification over coarse-grained propositions. In other words, a schematist may accept that we need to know *some* proposition to count as knowing modus ponens, but deny that this knowledge is absolutely general. We can make the point most clearly by using a logically true schema (i.e., a zero-premiss rule) such as ' $\phi \vee \neg\phi$ ', to avoid the form problem. Unlike the open-ended first-order schema ' $\alpha = \alpha$ ', which requires potential expansion to languages including terms for *any object whatsoever*, a schema such as ' $\phi \vee \neg\phi$ ' need only allow for expansions including sentences expressing *any proposition whatsoever*—and if propositions are coarse-grained, à la Stalnaker (1984), there may be many fewer propositions than objects.

It should be clear that although the concessive absolutist allows that the schematist can accommodate knowledge of inference rules, they agree with their strident counterpart that the schematist is stuck with an epistemological problem when it comes to schemata such as ' $\alpha = \alpha$ ' which are designed to mimic absolutely unrestricted quantification.

8 | CONCLUSION

Debate as to the coherence of absolutely unrestricted quantification has, until now, been almost entirely logically and metaphysically focused. We argue for a third axis of evaluation: the epistemological. We contend that the attempt, on the part of relativists, to eschew absolutely unrestricted quantification while using open-ended schemata to express absolutely general claims is epistemologically untenable.

REFERENCES

- Besson, C. (2018). Norms, reasons, and reasoning: A guide through Lewis Carroll's regress argument. In D. Star (Ed.), *The Oxford handbook of reasons and normativity* (pp. 504–28). Oxford: Oxford University Press.
- Carroll, L. (1895). What the tortoise said to achilles. *Mind*, 4, 278–80.
- Cartwright, R. (1994). Speaking of everything. *Noûs*, 28, 1–20. <https://doi.org/10.2307/221591>
- Chappell, S.-G. (2012). Varieties of knowledge in Plato and Aristotle. *Topoi*, 31, 175–90. <https://doi.org/10.1007/s11245-012-9125-z>
- Corcoran, J. (2006). Schemata: The concept of schema in the history of logic. *Bulletin of Symbolic Logic*, 12, 219–40. <https://doi.org/10.2178/bsl/1146620060>
- Dummett, M. (1991). *The logical basis of metaphysics*. Cambridge, Mass.: Harvard University Press.
- Dummett, M. (1995). *Frege: Philosophy of mathematics*. Cambridge, Mass.: Harvard University Press.
- Evans, G. (1982). *The varieties of reference*. Oxford: Clarendon Press.
- Fara, M. (2005). Dispositions and habituals. *Noûs*, 39, 43–82. <https://doi.org/10.1111/j.0029-4624.2005.00493.x>
- Feferman, S. (1991). Reflecting on incompleteness. *Journal of Symbolic Logic*, 56, 1–49. <https://doi.org/10.2307/2274902>
- Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48, 15–34. <https://doi.org/10.1007/BF00372404>

- Fricker, M. (2007). *Epistemic injustice: Power and the limits of knowing*. Oxford: Oxford University Press.
- Glanzberg, M. (2004). Quantification and realism. *Philosophy and Phenomenological Research*, 69, 541–72. <https://doi.org/10.1111/j.1933-1592.2004.tb00518.x>
- Hawthorne, J., & Manley, D. (2012). *The reference book*. Oxford: Oxford University Press.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32, 127–36. <https://doi.org/10.2307/2960077>
- Lackey, J. (2007). Norms of assertion. *Noûs*, 41, 594–626. <https://doi.org/10.1111/j.1468-0068.2007.00664.x>
- Lackey, J. (2009). *Learning from words: Testimony as a source of knowledge*. Oxford: Oxford University Press.
- Lavine, S. (1994). *Understanding the infinite*. Cambridge, Mass.: Harvard University Press.
- Lavine, S. (2006). Something about everything: Universal quantification in the universal sense of universal quantification. In G. Uzquiano & A. Rayo (Eds.), *Absolute generality* (pp. 98–148). Oxford: Oxford University Press.
- Lewis, D. (1991). *Parts of classes*. Oxford: Blackwell.
- Lewis, D. (1997). Finkish dispositions. *Philosophical Quarterly*, 47, 143–58. <https://doi.org/10.1111/1467-9213.00052>
- Martin, C. B. (1994). Dispositions and conditionals. *Philosophical Quarterly*, 44, 1–8. <https://doi.org/10.2307/2220143>
- Martin, C. B., & Heil, J. (1998). Rules and powers. *Philosophical Perspectives*, 12, 238–312. <https://doi.org/10.1111/0029-4624.32.s12.13>
- McGee, V. (1985). A counterexample to modus ponens. *Journal of Philosophy*, 82, 462–71. <https://doi.org/10.2307/2026276>
- Noë, A. (2005). Against intellectualism. *Analysis*, 65, 278–90. <https://doi.org/10.1111/j.1467-8284.2005.00567.x>
- Phillie, P. (2007). Carroll's regress and the epistemology of logic. *Philosophical Studies*, 134, 183–210. <https://doi.org/10.1007/s11098-005-7269-z>
- Pritchard, D. (2009). Safety-based epistemology: Whither now? *Journal of Philosophical Research*, 34, 33–45. <https://doi.org/10.5840/jpr/2009/2>
- Pryor, J. (2013). Problems for credulism. In C. Tucker (Ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism* (pp. 89–131). Oxford: Oxford University Press.
- Quine, W. V. O. (1937). New foundations for mathematical logic. *American Mathematical Monthly*, 44, 70–80. <https://doi.org/10.1080/00029890.1937.11987928>
- Rose, D., & Schaffer, J. (2013). Knowledge entails dispositional belief. *Philosophical Studies*, 166, 19–50. <https://doi.org/10.1007/s11098-012-0052-z>
- Ryle, G. (1945). Knowing how and knowing that. *Proceedings of the Aristotelian Society*, 46, 1–16.
- Smith, M. (2017). *Between probability and certainty: What justifies belief*. Oxford: Oxford University Press.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, Mass.: MIT Press.
- Stanley, J. (2011). *Know how*. Oxford: Oxford University Press.
- Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy*, 98, 411–44. <https://doi.org/10.2307/2678403>
- Studd, J. P. (2019). *Everything, more or less: A defence of generality relativism*. Oxford: Oxford University Press.
- Uzquiano, G., & Rayo, A. (Eds.) (2006). *Absolute generality*. Oxford: Oxford University Press.
- Whitehead, A. N., & Russell, B. (1927). *Principia mathematica* (2nd ed.). Cambridge: Cambridge University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (2003). Everything. *Philosophical Perspectives*, 17, 415–65. <https://doi.org/10.1111/j.1520-8583.2003.00017.x>
- Williamson, T. (2007). *The philosophy of philosophy*. Oxford: Blackwell.
- Williamson, T. (2009). Probability and danger. *Amherst Lecture in Philosophy*, 4, 1–35.
- Williamson, T. (2013). *Modal logic as metaphysics*. Oxford: Oxford University Press.