Non-citable pre-print. Forthcoming in Canadian Journal of Philosophy.

Anti-Exceptionalism about Logic and the Burden of Explanation

Ben Martin (University of Bergen)[†]

Abstract. Considerable attention recently has been paid to anti-exceptionalism about logic (AEL), the thesis that logic is more similar to the sciences in important respects than traditionally thought. One of AEL's prominent claims is that logic's *methodology* is similar to that of the recognised sciences, with part of this proposal being that logics provide explanations in some sense. However, insufficient attention has been given to what this proposal amounts to, and the challenges that arise in providing an account of explanations in logic. This paper clarifies these challenges, and shows how the *practice-based approach* is best placed to meet them.

Keywords: Logical explanations · Epistemology of logic · Methodology of logic · Anti-exceptionalism about logic · Logical predictivism · Practice-based approach

1. Anti-Exceptionalism About Logic

With the choice between numerous competing logics now facing logicians, it's no surprise that there has been an increased interest in logic's epistemology. One widely discussed position in the literature is anti-exceptionalism about logic (AEL), a cluster of theses proposing that logic is similar in important respects to the recognised sciences (Hjortland 2017; Williamson 2017). One of the prevalent claims made by advocates of AEL is that logic's *methodology* is similar to that of the sciences (Martin & Hjortland 2021 & forthcoming; Priest 2016). This proposal, however, has recently come under criticism for lacking both the necessary detail to be properly assessed (Martin & Hjortland 2021; Rossberg & Shapiro 2021), and sufficient positive evidence in its support (Martin 2021). This has led to the recognition of the need for proponents of AEL to both specify *which properties* of scientific methodology are mirrored within logic (and how), and provide evidence for these proposed similarities.

Some initial progress has been made, through the work of these same authors. Martin (2021) has argued that the most reliable route to providing evidence for an account of logic's methodology (anti-exceptionalist or not) is by looking in detail at logician's practice, using the so-called *practice-based approach*. Further, Martin & Hjortland (2021) have used the approach to provide a model of logic's methodology, *logical predictivism*, according to which predictive success plays a significant role within theory choice in logic, as it does in many of the recognised sciences. This work, however, explicitly leaves several important questions unanswered, such as which further theoretical virtues play a role in logical theory choice. One such open question, which this paper focuses on, is the extent to which logics (and logicians) *provide explanations*?

It has not been uncommon for those proposing anti-exceptionalist accounts of logic's methodology to note that logics could provide explanations of a kind (Payette & Wyatt 2018; Priest 2016; Williamson 2017), with Martin & Hjortland (2021) even hypothesising *explanatory power* as a further *desideratum* of logical theories. Yet, with the exception of Payette & Wyatt (2018), none have gone on to detail in what sense logics (and logicians) provide explanations, nor provide ample reasons for thinking that they indeed do. Thus, as

[†] Contact details: Benjamin.martin@uib.no

with the anti-exceptionalist proposal in general, more needs to be said about the specifics of what's being proposed with regards to explanation, and the proposal's rationale.

Our goal here is to advance the debate, by: (i) clarifying the burden that the methodological anti-exceptionalist takes on in demonstrating that logics explain; (ii) demonstrating how the only currently available account of logical explanation—in Payette & Wyatt (2018)—fails to meet this burden; and, (iii) outlining ultimately how the anti-exceptionalist should go about meeting it, using the aforementioned *practice-based approach*.

The paper runs as follows. Section 2 outlines methodological AEL, and clarifies the anti-exceptionalist's burden regarding explanations in logic. Section 3 introduces the *practice-based approach*, and argues for why it's the best method to broach the topic of logical explanations. The case for using the *approach* is then strengthened in Sections 4 & 5, by showing that it doesn't fall foul of the pitfalls of Payette & Wyatt's (2018) recent discussion. We finish in Section 6 by emphasising the dangers of drawing too close a connection between explanations in logic and those in other (scientific) fields.

2. Methodological Anti-Exceptionalism and Logical Explanations

2.1 Methodological AEL

AEL is the view that logic is not special in one or several of the ways in which it has been traditionally considered to be, such as its *apriority*, analyticity, and normativity (Hjortland 2017; Martin & Hjortland 2021). One of the most important, and prevalent, anti-exceptionalist views is that logical evidence is not immediate and foundational in the way that those who have traditionally appealed to rational insight (BonJour 1998) and epistemic analyticity (Ayer 1936) suggest. Rather, the method through which we come to be justified in holding logical propositions, and theories, is similar to that of the sciences:

Methodological AEL: The *methodology* of logic is similar to that of the recognised sciences.

Yet, whilst one of the significant motivations for methodological AEL is an attempt to de-mystify logic's epistemology, and avoid the concerns associated with appeals to intuition and epistemic analyticity (Martin 2021), the proposal has come under criticism for its own shortcomings.

Firstly, it has been criticised for its lack of detail in elucidating what the resulting methodology amounts to, which subsequently makes the thesis difficult to assess (Rossberg & Shapiro 2021). For instance, there are clearly prevalent components of scientific methodologies which are not mirrored within logic (Martin & Hjortland 2021 & forthcoming), such as the development of techniques and apparatus in order to measure important constants with increasing levels of accuracy and precision, particularly in the physical sciences (Tal 2020). Logic, in comparison, has no obvious interest in measuring constants (though, like many disciplines, its practitioners develop new techniques for various purposes). Further, statistical methods are rife across the natural and social sciences; yet, whilst these methods *make use of logic*, logicians themselves make no use of statistical methods in developing or assessing logics. Thus, more needs to be said about *which* purported features of scientific methodology occur within logic.

These concerns persist even when we consider one of the more detailed (and popular) accounts of how logic's methodology is similar to that of the sciences—*logical abductivism*—according to which logic's

theories are justified and selected for in accordance with an abductive method, just like (purportedly) scientific theories are (Priest 2014; Williamson 2017). Such accounts often lack detail on how to make sense of logics "fitting" data, or how to understand particular theoretical virtues (such as simplicity) in the context of logic (Martin & Hjortland 2021).

Further, as exemplified by presentations of logical abductivism, methodological AEL is also often lacking significant positive evidence in its support. Whilst some advocates of abductivism are explicitly motivated by the weaknesses of past attempts to explain logic's epistemology in terms of rational intuition or epistemic analyticity (Sher 2016; Williamson 2017), it's striking how little positive evidence has been advanced for logical abductivism itself (Martin 2021). In some of the papers often cited as *defences* of abductivism (e.g. Priest 2014), what we actually find are models for how logical theory choice *could* function by abductivist means, not how it *does* or *should*. Yet, of course, the fact that competing accounts of logic's epistemology suffer significant weaknesses does not itself provide positive support for logical abductivism, given that abductivism does not exhaust our other options. Indeed, non-foundationalist epistemologies of logic (including abductivism) are well-known to suffer their own problems, such as the *background logic problem* (discussed in Section 4), and this could well give one reason to think that *any* cognitivist picture of logic's epistemology is implausible (Wright 1986 & 2018). Consequently, in order to show that abductivism, or any other anti-exceptionalist methodology of logic, is indeed superior to competitors, greater positive evidence must be offered in its support.

This is the backdrop against which discussions of methodological AEL must proceed: both greater clarity over its proposals, and independent evidence for them, are needed. While progress has been made on this score (Martin 2021; Martin & Hjortland 2021), these works leave several questions about logic's methodology open. Our focus here is on one of these proposals: that logics (and logicians) are *engaged in providing explanations*. Particularly, in keeping with the challenge facing methodological AEL, we are interested in: (i) What sense can be made of the proposal? and, (ii) How can we go about testing it?

2.2 Clarifying the Task

Firstly, we need to specify the type of explanations which the anti-exceptionalist is interested in. After all, there are certain explanatory activities which uncontentiously occur within logic, and thus cannot plausibly be the target of her proposal. For instance, logic textbooks commonly provide *pedagogical explanations* of appropriate disciplinary standards and field-specific tools for students, such as how to prove a theorem within a logic L proof-theoretically, and which metalogical results the community cares about. While such pedagogical explanations clearly occur within logic, they are also a part of many other enterprises, scientific or not, and thus cannot reasonably be the bone of contention here.

Compare this with two other types of explanation—*intra-systemic* and *extra-systematic explanations*—distinctive of scientific methodologies. *Intra*-systemic explanations provide information about the field's own artefacts, whether they be theories, models or logical systems, without any concern given to the applicability of these artefacts to phenomena external to them. Here the artefacts are studied as objects of interest in and of themselves, and researchers are interested in explaining their *intrinsic* properties, or relation to other artefacts. Providing such *intra*-systemic explanations constitute an important part of the mathematical sciences, and fields that make heavy use of mathematical modelling. For instance, while epidemiological models are often *applied* to real-life diseases, mathematical epidemiologists are also

interested in explaining the properties of particular epidemic models, and comparing them with other available models, *without* reference to the models' potential application to actual diseases or predictive success (e.g., Tuckwell & Williams 2006).

In comparison, *extra-systematic* explanations are focused on the applicability of a field's models to independent phenomena, and in particular their ability to explain features of these phenomena. Thus, in these cases, rather than being interested in explaining the intrinsic qualities *of the model*, one is interested in *using the model* to fruitfully explain some external phenomenon, such as the spread of diseases in the case of epidemiology. *Extra*-systematic explanations are found across all areas of the natural and social sciences: stellar astrophysicists are engaged in the explanation of the evolution of stars, molecular biologists in explaining the various systems of cells, and criminologists the social causes of crime. While not every researcher will be engaged in providing extra-systemic explanations—some will be more interested in improving experimental or statistical techniques—they constitute a significant part of the sciences.

In the case of logic, it's quite clear that logicians are actively engaged in providing *intra*-systemic explanations. Logics can be, and are, produced and studied for their own sake, without reference given to their potential applications. Logicians are interested in producing systems with certain intrinsic properties, proving that these systems have these properties, and subsequently explaining their relationship to other existent systems. One of the more simplistic and straightforward examples of such *intra*-systemic explanations within logic is the explanation of why a particular formula is a theorem of a logic given its semantics. Logicians are able to explain, for instance, why the law of excluded middle is valid in classical logic but not in the gappy K3. These are explanations of what is *valid in the particular logic L*, given *L*'s semantics, and how this compares with other logics, without reference to the applicability of these logics to some independent phenomena. Thus, just as with the (other) mathematical sciences, logic is engaged in providing *intra*-systemic explanations. Consequently, if *intra*-systemic explanations were the subject of the anti-exceptionalist's proposal, there would be little to get animated about.

In contrast, what is not so clear is the extent to which logics provide *extra*-systemic explanations of some given external phenomenon. Demonstrating that they do requires not only highlighting suitable target phenomena which logics can be applied to, but that our logics have a suitable relationship to these phenomena worthy of being called "explanatory". Neither claim is obviously true, and as we shall see both have been called into question.

Yet, it is seeming this proposal that logics can, and do, provide *extra-systematic* explanations which is the substance of the anti-exceptionalist's thesis that *logics explain*. Specifically, while there are various phenomena which logicians appear to be interested in applying their systems to, including belief revision (Hansson 2017) and grammatical structures (Dalrymple 2001), anti-exceptionalists are particularly focused on demonstrating that logics can provide *extra-systematic* explanations of the putatively philosophically important phenomenon, *validity*:

> Anti-exceptionalist views of logical theory choice are really about which logic to adopt as our most basic canon of legitimate deductive implication. (Woods 2019: 320)

> A decent logical theory is no mere laundry list of which inferences are valid/invalid, but also provides an explanation of these facts. (Priest 2016: 353)

> [W]hat we are looking to explain is the validity/invalidity of some argument consisting of (fully interpreted) natural language sentences. (Payette & Wyatt 2018: 159)

While foundationalist accounts of logic's epistemology propose that we have direct access to knowledge of validity, whether in terms of rational insight or epistemic analyticity (Martin 2021), methodological anti-exceptionalists propose we gain knowledge of validity by constructing theories, which are at least partially evidenced by their ability to *explain* the target phenomenon. Consequently, it is logics' ability to provide extra-systemic explanations of validity which is the substance of the anti-exceptionalist's claim that logics explain:

> *Validity Explanation* (**VE**): Logics provide *extra-systemic explanations* of validity.

Clearly, **VE** is far from trivial. There are at least two reasons why one might reject **VE**, both with prominent advocates. Firstly, one could deny there are any such facts about validity for logics to explain. Those who propose that our logical theories serve an expressivist function, such as Resnik (1999), are good examples of those who would reject **VE** for this reason. Logics merely *express* our own predilections, rather than *explaining* why a particular argument is valid.

Secondly, one could admit that although there are facts about validity, our logics fail to offer any explanation of this phenomenon, as they provide us with no further understanding of an argument's validity beyond what we already know. A varied group have held this view, including historically Descartes and Mill, and in the contemporary literature Szabó (2012). While their exact rationales differ slightly due to their underlying philosophical commitments, for both Descartes (Gaukroger 1989) and Mill (Godden 2017) logics simply provide us with schematic generalisations of instances of arguments that we already recognise to be valid (due to rational insight or experience). Thus, logics provide us with a neat means to *represent* validity relations between propositions, but fall short of *explaining validity*. For both, the schematic generalisations that logics offer are no more insightful and explanatory than the generalisation that "All swans are white" is of why a particular swan is white.

These challenges to **VE** have consequences for the anti-exceptionalist. Whereas the philosopher of science tasked with elucidating the nature of particular extra-systemic explanations can generally rely upon the assumption that the field is engaged in such activities, and further will possess paradigm examples to work from, the anti-exceptionalist has no such privilege. There is no justified presumption in favour of the existence of explanations of validity in logic—it is exactly the truth of this claim which is in question. Thus, she is obliged to *both* (i) supply us with reasons to think that logics do indeed provide *explanations* of validity, and (ii) elucidate *how* logics provide such explanations.

In the remainder of this paper, we'll show how the *practice-based approach* is best placed to meet this burden. While the next section outlines the approach, the proceeding sections show how the approach avoids the pitfalls of a recent discussion of logical explanations.

3. The Practice-Based Approach

According to the *practice-based approach*, the most reliable means to come to know the aims, objectives and methodology of logic is through the actions of its practitioners (Martin 2021 & forthcoming). As with

studies into the methodologies of the sciences, in order to understand what constitutes *logic* and how its theories are evidenced, we are best off engaging not in conceptual analysis or deliberating over the *nature of logic itself* (whatever that is), but looking at how logic is actually practiced. We must form hypotheses about its aims and methodologies, and test these expectations against the activities of practitioners in the field.

The justification for this proposal comprises two parts. Firstly, like all fields of enquiry, logic is a social construction (though its objects of study definitely need not be). This means that its interests, aims and techniques are the result of the collective actions of its practitioners, and are not predetermined by some other factor. Further, these interests, aims and techniques can change significantly over time. This has two immediate consequences for anyone interested in understanding logic's aims and methodology: Firstly, one cannot hope to come to a detailed and reliable account of the field's goals and methodology without engaging with how it's actually practiced, as it is these collective practices which *constitute* its goals and methodology. Thus, there's no direct route to this understanding through conceptual analysis or thought experiments. Secondly, one cannot appreciate logic's current goals and methodology by appealing wholly to those of its past. While there may be continuities between a field's past and present goals and methodology, it's difficult to establish these without looking at current practice.

Importantly, recognition that the *field* of logic is a social construction whose goals and methodologies change over time does *not* entail that its object of study or the knowledge it produces are social constructions. Rather, it is the field's *goals and methods* which are the results of its participants' collective actions. Further, this recognition should not call into question the field's rationality: theoretical decisions can still be wholly rational and improved over time. For any field of enquiry containing expertise, one can expect its aims to become progressively more precise, its techniques for realising these aims more reliable, and its theories relative to these aims more successful (Martini 2014). Theoretical progress, therefore, is completely compatible with recognising that research fields themselves are social constructions.

This brings us onto the second part of our justification for the *approach*. Being experts within the field, we can expect logicians to be the most reliable judges of the suitability of a particular research question, the most fruitful techniques in the given scenario, and what constitutes relevant evidence for possible answers to the question. It is exactly this which makes them experts (Martini 2014). Logicians are not unique in this regard. The exact same point holds for other experts in their field, including scientists and legal professionals. In virtue of being experts, we expect these individuals' professional practices to reliably exemplify the field's aims, methods and virtues. This is exactly why a considered account of scientific methodology (Bokulich 2011; De Regt 2017) or theory of jurisprudence (Dworkin 1986) would be found inadequate if it failed to make sense of a considerable proportion of the practice of experts in the field. Combined, these considerations demonstrate that logicians' practices are the most reliable guide we have to logic's methodology.

Before moving onto discuss how the *approach* can help us test **VE**, it's important we first remove two potential confusions over the *approach*. Firstly, while the *approach* argues that the practice of logicians as a whole are a reliable guide to logic's methodology, this does not mean that individual expert's practices are always reliable. Even experts make mistakes. For this reason, it's important one takes into account the activities of multiple experts, and the responses of their peers to these activities, in order to ensure one builds a *representative* picture of the relevant practice (Martin 2021).

Secondly, while experts' field-specific practices are a reliable guide to the field's aims and methodologies, this does not mean their *reflective views about the field* are a reliable guide to these features

of the field. Reflective views are not always representative of one's own practice. Good scientists do not necessarily make good philosophers of science, and excellent natural speakers of Spanish do not necessarily have accurate opinions about the rules constituting grammaticality in Spanish (this is why we have descriptive linguistics). We must, then, be conscious to distinguish between instances of what practitioners *actually do* in the relevant practice and what they *say about it*. While the former is a reliable guide to the workings of the field, the latter is not. Thus, just as the philosopher of science who is attempting to build an account of scientific methodology does not *ask* scientists what they think constitutes the scientific method, but rather observes their various professional practices in order to infer the norms underpinning the field's methodology. Similarly, the philosopher of logic must look at the research questions logicians ask, the techniques they use in answering these questions, and the types of evidence they appeal to in justifying their logics, to acquire an accurate picture of logic's methodology (Martin 2021).

In order to build our positive case for using the *practice-based approach* to assess **VE**, we'll take as our starting points the shortcoming of the only existent detailed discussion of logical explanations, by Payette & Wyatt [hereafter, P&W] (2018). P&W (2018) propose that logical explanations are best understood in terms of Andrea Woody's (2015) *functional perspective* of explanation. Their case for this claim comprises two stages. The first is to show that what they call "traditional" accounts of explanation, such as the deductive-nomological model, are wholly unsuitable for the anti-exceptionalist's purposes. In comparison, we're told, Woody's *functional perspective* falls foul of none of these concerns, and so can be used to characterise logical explanations. The second stage is to argue that if the anti-exceptionalist takes up Woody's *functional perspective*, she will be able to provide an account of logical explanation as similar to that in (at least) some sciences.

In what follows, we'll argue that P&W are mistaken on three counts: (i) They are too quick to dismiss the appropriateness of certain models of explanation for logic; (ii) Woody's *functional perspective* is insufficient on its own to meet the anti-exceptionalist's burden of justifying **VE**; and (iii) They draw too close of a connection between logical explanations and those in other sciences. In comparison, the *practice-based approach* allows us to avoid each of these problems, further highlighting the suitability of the approach to properly assess **VE**.

We begin with P&W's case for why "traditional" accounts of explanation are unsuitable to characterise logical explanation, why this is mistaken, and how it could distract the anti-exceptionalist from finding an adequate justification for **VE**.

4. Prejudging Models of Explanation

For P&W (2018: 159), "traditional" accounts of explanation are those which attempt to provide adequacy criteria for individual explanations. Examples include the deductive-nomological (**DN**), causal and mechanistic models. While each model suffers its own potential shortcomings (Woodward 2017), P&W (2018: 159-60) propose that a greater challenge faces their application to logical explanations, given that each model requires the presumption of certain rules of implication. In order to clarify this *circularity problem*, we'll focus as P&W do on the example of the **DN** model.

According to the **DN** model, explanations take the form of arguments, with the phenomenon to be explained as the conclusion, and the proposed *explanans* as the premises, of which there are two types: the set of propositions representing the initial conditions, and those representing the relevant laws. The

explanation is deemed successful if and only if the argument is sound and contains a law essential to the argument's validity. Applying the model to the case of logic, P&W (2018: 159) propose we understand the initial conditions to express how particular natural language arguments are formalized in the logic, and the laws to express the proposed logic's semantics, however conceived. Taken together, the initial conditions and laws either *entail* that the relevant vernacular argument is valid or invalid. If the explanatory argument turns out to be sound, then the explanation of the validity of the target argument is deemed successful.

Thus, applied to logical explanations, the **DN** model would require us to *presuppose certain logical laws* (rules of implication) in order to show that the logical laws (in conjunction with initial conditions) proposed as *explanantia* adequately explain the *explanandum*. It is this requirement which, according to P&W, makes the **DN** and other "traditional" models characteristically unsuitable to elucidate logical explanations. If logicians were to use the **DN** model, they would be forced to assume what they are seeking to establish—the truth of certain logical laws.[1] It's for this reason that P&W propose we look instead to Woody's *functional perspective* for an account of logical explanations, as they claim it does not require a logical relationship between the proposed *explanans* and *explanandum*.[2]

P&W are right that this apparent circularity constitutes a philosophical puzzle for anyone wishing to propose that logical explanations exemplify the **DN** model, and similarly for any model requiring deduction within the explanatory process. However, they are mistaken that this is reason enough for the anti-exceptionalist to abandon these models when accounting for logical explanations.

Firstly, the fundamental epistemological problem of having to rely upon logical rules of implication in order to support a logical theory, of which P&W's *circularity problem* is a version, would persist for the anti-exceptionalist *even if* she were to reject these "traditional" accounts of explanation. This wider epistemological problem, which is a well-known feature of non-foundationalist accounts of logic's epistemology (Sher 2016; Shapiro 2000), is known as the *centrality* (Wright 1986) or *background logic* (Martin 2021) *problem*.

According to the *background logic problem*, any epistemology of logic which proposes that we come to be justified in believing a logic *L* by appealing to non-immediate evidence *E* will come unstuck, as we will always need to appeal to logical rules in order to demonstrate that *E* is (in)consistent with *L*. In other words, the justificatory process requires making *deductive inferences*. However, of course, any logical rules relied upon in making such inferences will either need to be sanctioned by the logic *L* under consideration or not. If they are, then the logic is simply begging the question in relying upon the rule for its evidential support, and thus the putative justification offers no new evidence for the theory. Instead, if the theory precludes the validity of the deduction, then the theory undercuts its own possible justification. Either way, logics' justification cannot suitably rely upon non-immediate evidence, as it requires deductive inferences to be

[1] Underlying this worry is the assumption that using *rules of implication* to deduce results from the *logical laws* equates to presupposing the truth of the logical *laws* in question. Yet, rules of implication are not themselves theorems; though, they can be presented as such with an appropriate Deduction Theorem. Consequently, there's a concern that P&W are unjustifiably presuming we could not be warranted in inferring according to particular rules of implication (in certain circumstances) *without presuming the truth of the relevant logical laws* in the proposed *explanans*. We won't push this point any further here, however.

[2] Whether this is correct is doubtful, as we'll see in the next section.

made. In the same respect, according to P&W, logical explanations cannot rely upon a relationship between the *explanandum* and *explanans* that presupposes deductive inferences.

The upshot of this connection between the *circularity* and *background logic problems* is that P&W have overemphasised the impact the *circularity problem* should have on the anti-exceptionalist. Whatever happens to their account of logical explanation, the anti-exceptionalist will need to solve the *background logic problem*, given that a significant motivation for AEL is that there is no immediate *a priori* foundation for logical justification (Martin & Hjortland 2021). Instead, we must build up evidence for our logical theory, whether based upon linguistic judgments, results from mathematics, or the logico-semantic paradoxes, all of which rely upon *arguments being given* for why these data constitute evidence for the logic. Thus, refusing to use certain accounts of explanation because they fall foul of the *circularity problem* would be fruitless for the anti-exceptionalist unless she also had a solution to the *background logic problem*.

Yet, as the *circularity problem* is but an instance of this wider *background logic problem*, it's likely that a solution to the latter problem would also provide a solution to the former. After all, solving the *background logic problem* will ultimately require showing that logicians can be warranted in using certain rules (or, instances of those rules) to test the consistency of a logic with some given data, *even if those rules are contested*. Yet, if this can be shown, it would also demonstrate how logicians could be warranted in using certain rules (or, instances of them) to establish a relationship between some *explanantia* and *explananda* without begging the question.

Thus, using the *circularity problem* to motivate a particular account of logical explanation for the anti-exceptionalist is misguided, given that it's already well recognised that the anti-exceptionalist must ultimately provide a solution to the background logic problem. Accordingly, P&W's motivation for why the anti-exceptionalist ought to prefer using Woody's *functional perspective* over traditional models is undercut.

There is a deeper problem here, however, with P&W's criticism of "traditional" models on the basis of the *circularity problem*, which speaks in favour of the *practice-based approach*. It's likely that attempting to motivate an account of logical explanation via the *circularity problem* will unduly prejudice against otherwise plausible accounts of logical explanation, and thus distract the anti-exceptionalist from the best possible routes to justifying **VE**. The anti-exceptionalist's task is to show that logicians are engaged in providing extra-systemic explanations, and to elucidate the nature of these explanations. There is no requirement that the resulting explanatory models, if there are any, are free from philosophical puzzles. We do not require this of scientific methodology, and there is no need to require it of logic's methodology. Indeed, rejecting out of hand the viability of these types of explanation within logic could put us in a tricky situation, for how should we proceed if we *do* find logicians using the **DN** model to support their theories? We can hardly reply that logical explanations *cannot* function this way, any more than we could for explanatory practices in the sciences. If we find logicians using the **DN** model, our goal is to *find solutions* to the philosophical problems these models of explanation produce, not to deny their status as viable explanations! A philosophically puzzling feature of a theoretical practice does not itself constitute reason to doubt its existence. Consequently, we should be wary of unduly prejudicing against possible accounts of logical explanation, by outright rejecting their viability. Fortunately, the *practice-based approach* helps us avoid this very pitfall, by recognising it is logicians' practices which ultimately dictate the form(s) that logical explanations take.

How though do we go about analysing explanatory practices within logic? To even begin *analysing* these practices, we must first be capable of *identifying* them. Given that the existence of extra-systemic

explanations within logic is itself moot, this raises certain challenges. To understand them, and how the *practice-based approach* can help solve them, first we will consider how P&W's appeal to the *functional perspective* fails in this regard.

5. Substantiating Logical Explanations

The functional perspective itself is an attempt to re-orientate discussions about scientific explanation back towards scientific practice, and take seriously what scientists within various fields take to be explanatory. Rather than relying upon intuitions over whether out of context cases constitute explanations or not, which are unreliable guides to what's explanatory in the sciences, we should look at what scientists treat as explanatory in their practice (Woody 2015: 81).

Further, as we have significant evidence that scientific explanatory practice is not homogenous (Ruphy 2017: Ch. 3), it's important that our analysis does not focus on singular explanations. This could lead to hasty generalisations from these particular cases and easily slide into "unwarranted essentialism about the nature of explanations across the sciences," declaring "whole categories of explanations tendered by practitioners illegitimate or inadequate" (Woody 2015: 80).[3] To circumvent this concern, Woody (2015: 79) recommends our analysis of explanatory practices moves from the traditional question of the *adequacy conditions for individual explanations* to the question of what function(s) explanations play within the sciences.

It is for this reason that P&W (2018: 160) propose anti-exceptionalists are best served using the *functional perspective* in elucidating logical explanations. Putatively, it allows the anti-exceptionalist to avoid the *circularity problem* by concentrating on the *function* of logical explanations, rather than attempting to account for an individual explanation's success in terms of some logical relationship between an *explanandum* and *explanans*. Unfortunately, however, contrary to what P&W suggest, the approach is both ill-equipped on its own to aid the anti-exceptionalist in justifying **VE**, and insufficient to avoid P&W's circularity problem. We'll deal with this latter point first.

By requiring us to re-focus our attention to the *function* of explanatory practices within the sciences, the perspective does not somehow dissolve the question of the properties of individual explanations within the sciences (or logic). Rather, it simply sidesteps the question for the time being, and asks us to consider the wider role these explanatory practices play in the field. Woody (2015: 81) herself is clear that the *perspective* does not pre-judge the nature of these individual explanations. There are *still* instances of explanation within each field, and they will *still* have their own properties. For all we know, some logical explanations *could* conform to the **DN** model. Thus, we do not solve the circularity problem for logical explanations by simply moving our perspective onto the *function* of explanatory practices within logic; we sidestep the issue. If, ultimately, the anti-exceptionalist wished to show that extra-systemic explanations in logic did not fall foul of the *circularity problem*, she would still be required to look at the form individual explanations take within logic. Thus, the *functional perspective* is not the silver bullet P&W propose.

Even more concerning, however, is the fact that the *functional perspective* is not capable unaided of helping the anti-exceptionalist justify **VE**. The anti-exceptionalist is in the business of *both* establishing that logicians provide extra-systemic explanations, and then subsequently elucidating their features. Yet, the functional perspective *presupposes* the existence of instances of explanations in a field, from which we can

[3] Note the irony here, that Woody is warning against exactly what P&W were shown to be guilty of in Section 4.

then draw implications about their functions. Without initial agreement on instances as cases of explanatory practice, there are no data for the functional perspective to get off the ground (Woody 2015: 81). While this is fine in the case of the established sciences, where it's uncontroversial that practitioners are engaged in providing such explanations and we can point to paradigm instances, in the case of logic the existence of such explanations is itself moot. Presuming the existence of such explanations within logic will, ultimately, do the anti-exceptionalist no good. She needs to find the means to independently identify instances of logical explanations, and this cannot be achieved through the functional perspective itself.

Prima facie, the same quandary faces any use of the *practice-based approach* to support **VE**, given that the *approach* requires us to point to instances of such explanatory power in order to show that logicians provide extra-systemic explanations. Yet, of course, to even recognise instances of practice *as cases of extra-systemic explanation*, one must already have a good sense of what would constitute such an explanation in logic. Without this, one is like a bird-watcher attempting to spot a new species of sparrow without an indication of what it would look like, or indeed whether it even exists.

Commonly, discussions of scientific explanations are conducted either from a *bottom-up* or *top-down* perspective (Braillard & Malaterre 2015). Unfortunately, however, neither are suitable for the anti-exceptionalist's purposes. The *bottom-up* approach begins with paradigm instances of explanatory practice within the field and extrapolates from these to an account of explanation in the research area. This is generally how accounts of explanation in the sciences now proceed (Bokulich 2011; Fagan 2015), given the appreciation that explanatory norms can be field specific. This option simply isn't available for the anti-exceptionalist, however, given that it presupposes the existence of paradigm cases of explanations in the field.

The alternative, a *top-down* approach, would be to presume a particular model of explanation and demonstrate that instances of logical practice fit *that* model. This approach is associated with earlier attempts to provide a *universal* account of explanation, such as the **DN** model or Salmon's (1971) *statistical relevance* model. Yet, this option faces its own complications. In particular, there is no universal model of explanation which can successfully capture all scientific explanations. Instead, there are various types of explanatory models used by scientists across disciplinary boundaries. This is one of the reasons behind the prevalence of the *bottom-up* approach in the contemporary literature.

Of course, there's nothing to stop the anti-exceptionalist from simply trying each of these available models out, and searching for instances of logical practice which fit. However, given that explanatory norms differ from subject-to-subject (Woody 2015), it's unclear that, even if **VE** were true, one of the existent models of explanation would neatly fit logical explanatory practice. Further, as Woody (2015: 80) herself has stressed, explanatory practice is always context-dependent. No scientific (or logical) practices are explanatory purely in virtue of their intrinsic properties, but instead are dependent upon the particular aims and norms of the community. This means that simply identifying an instance of practice which fits a model of explanation does not entail that this practice is in fact *explanatory*. To fully understand what constitutes an explanation within a field, one has to appreciate its underlying goals and subject matter.

This final point gives us a clue as to how the anti-exceptionalist might meet her burden without simply presupposing the existence of extra-systemic explanations in logic. Her task is to show that certain practices within logic *deserve the honorific of being explanatory*, contra the doubts raised in Section 2. One way to achieve this is to point to the similarity between certain practices found within logic and those in other fields that we consider to be explanatory in these latter contexts. Given that explanatory norms are very much impacted by a field's goals and subject matter, being able to draw these connections will itself require having

insight into the aims and subject matter of logic, and of the comparative fields. It is unlikely, for example, that the anti-exceptionalist will find much joy in drawing a connection with probabilistic explanations exemplified in medicine, given the lack of the use of probabilistic tools in assessing logics.

The most promising approach to take, therefore, is to first acquire a good understanding of logic's aims and methods on the basis of logicians' practices. From this independently evidenced account of the field, one can then attempt to draw connections to other fields of enquiry that share certain of these aims and methods, with the ultimate goal of highlighting practices within these fields that are: (i) considered explanatory, and (ii) analogous to practices within logic. It is through establishing this connection to recognised explanatory practices in associated fields that the anti-exceptionalist will be able to provide a strong case for why practices within logic deserve the honorific of being explanatory.

In order for such an *argument by analogy* for **VE** to succeed, one requires three components. Firstly, one must possess an well-evidenced account of logic's aims and methodology prior to drawing any such analogy, using the *practice-based approach*. Secondly, one must have an informed account of explanatory practices in other fields, in order to draw the suitable analogies. Finally, one must have an argument for why similarities between these explanatory practices and those found in logic suffice for concluding that logic provides extra-systemic explanations.

While our aim here is not to provide a detailed defence and elucidation of extra-systemic explanations in logic, it will be instructive to briefly outline how such an argument for **VE** could proceed. Firstly, as we've noted, one needs a prior understanding of (a portion of) logic's aims and methods via the practice-based approach. For the sake of illustration here, we'll use a proposal recently defended using the *approach*, called *logical predictivism* (Martin & Hjortland 2021).

According to *logical predictivism*, one important aim of logical theories is to provide an account of validity, conceived as a property of arguments. In such cases, logics are justified through a combination of their predictive success, explanatory power and compatibility with other well-evidenced commitments. Importantly, so that logical theories can produce predictions to be tested against suitable data, these theories are not conceived of as simply a set of valid rules of inferences or theorems, but as a cluster of definitions, laws and representation rules that provide the underlying semantics and syntax of the theory, as well as specifying how it connects to the extra-systemic phenomenon. Here's a toy example of classical propositional logic under such an account:

Definition 1: Let $\neg\phi$ be Boolean negation.
Definition 2: Let $\phi \wedge \psi$ be Boolean conjunction.
Representation Rule 1: $\ulcorner$ not $\phi \urcorner = \ulcorner \neg\phi \urcorner$.
Representation Rule 2: $\ulcorner \phi$ and $\psi \urcorner = \ulcorner \phi \wedge \psi \urcorner$.
Law 1: For every valuation, all sentences are either true or false, and not both.
Law 2: An argument is valid iff, for every valuation *v*, if every premise is true in *v*, the conclusion is true in *v*.[4]

[4] We're passing over many of the nuances here, such as how to deal with hypothetical arguments; see Martin & Hjortland (2021) for details.

These theories (putatively) have three properties which are interesting for our purposes. Firstly, they include *idealisations*, in the form of their definitions and representation rules. Everyone accepts, after all, that not every use of "not" in English is equivalent to a truth-functional negation (Horn 1989). Secondly, they can include *fictions*, such as when theories appeal to (im)possible worlds in their semantics. Lastly, the theories *specify the conditions under which* arguments are (in)valid, and thus elucidate why particular arguments are valid and others invalid.

Once one has some understanding of logic's aims and methodologies, attention then moves to drawing suitable connections between practices in logic and recognised cases of explanatory practices in the sciences. Fortunately, philosophers of science have done much of the hard work for us here, identifying multiple forms of explanation across the sciences. Our attention then must move to looking for potential similarities between these explanatory practices and those of logic. In the case of *logical predictivism*, one such promising line of enquiry is to point to the similarity between logical theories so conceived and instances of *model explanations* that one often finds in the sciences (Bokulich 2012). Two features of these models are worthy of note.

Firstly, just as with *logical predictivism*'s account of logical theories, scientific models readily contain idealisations, abstractions, and fictions, which provide the models with the theoretical virtues they are prized for (Bokulich 2011). Secondly, according to an increasingly prominent account of what makes these models *explanatory*, they are capable of specifying how changes to elements of the model (the *explanans*) would result in changes to the *explanandum* (Bokulich 2011 & 2012), building on Woodward's (2003) *counterfactual* account of explanation. Thus, the models are explanatory in virtue of being able to capture patterns of counterfactual dependence in the target phenomenon, allowing us to answer a range of *what-if-things-had-been-different* questions about the phenomenon *using the model*.

In order to successfully build a case for **VE** based upon *logical predictivism*, and the proposed similarity between practices within logic and these model explanations in the sciences, one would need to successfully argue that the pertinent features which make these models explanatory in the sciences also obtain in logics. For instance, that just as with scientific models, logics (the *explanans*) are able to specify how changes to parameters within the theory, such as the logical form of a given natural-language argument, would result in changes to the *explanandum* (namely, the (in)validity of the arguments).

As we have emphasised, our example here is merely illustrative. There are various other possible accounts of logic's methodology and scientific explanations, suitably informed by practice, which one could use in attempting to draw the relevant connections. However, the example serves to clarify how such an argument by analogy could support **VE**, and highlights a future line of enquiry for advocates of AEL.

6. Drawing Conclusions from Practice

We end our discussion of **VE** by emphasising two important features of any *argument from analogy* for **VE**, and the potential dangers of failing to respect them. Doing so should not only be instructive for future attempts to establish **VE**, but highlight several weaknesses of existent proposals.

Firstly, in order for the argument to be successful, one requires a well-evidenced appreciation of logic's aims and methodology *prior to* drawing a connection with explanatory practices in other fields. If, instead, one simply begins with a certain picture of scientific methodology, and an intention to draw connections between scientific and logical methodology, one increases the risk of overemphasising any points of similarity

and thereby distorting logic's methodology. One needs to start from a solid base of appreciating the realities of logic's methodologies.

Secondly, in being an argument from analogy, the argument is not intended to show that explanations in logic are *identical* to those in other fields. Rather, it only serves to substantiate the claim that logics are engaged in providing extra-systemic explanations, by appealing to pertinent similarities with explanatory practices in other fields. For all we know, explanations across fields may hold the status of being explanatory in virtue of certain family relations. Consequently, once an *argument from analogy* has provided us with evidence for the occurrence of extra-systemic explanations in logic, the precise features of these explanations are then a matter to be decided by logic's practices, not those of another field. In other words, it's paramount we do not confuse **VE** with a stronger principle,

> *Equivalent Explanation* (**EE**): Logics provide extra-systemic explanations of validity, which share all of the pertinent features of extra-systemic explanations in field *F*.

and fall into the trap of presuming AEL is obliged to defend this stronger principle **EE** in virtue of endorsing **VE**.[5]

The importance of respecting both of these features can be demonstrated through a brief consideration of P&W's (2018) own account of logical explanations, which draws a close relation between logical explanations and those in the physical sciences. Using Woody's (2015) own functional analysis of the ideal gas law in chemistry as a starting point, P&W (2018: 162) argue that logical explanations play a similar function to those in chemistry.

According to Woody (2015: 82-3), while the ideal gas law is taught in university-level courses and used by practitioners, it's well-known to be empirically inadequate. Particularly, the law fails to take into account the size of the molecules in the gas and their interaction. Thus, depending upon whether a *ceteris paribus* condition is added to account for these extraneous factors, it ends up being either straight up false, or inapplicable to actual non-idealised gases. Despite this, the law still serves pedagogical and scientific functions, each of which inform us about the function of explanations within chemistry.

In particular, the law acts as a visual prompt, giving students a model of gases as constituted of uniform compact particles large distances apart which exert little force upon one another. Secondly, it specifies important properties which the student ought to pay attention to in gases, acting as "inferential scaffolding for the treatment of all gases" (Woody 2015: 82). Further, it acts as a barometer with which to judge the actual behaviour of gases against, with the resulting departure functions being important theoretical results in themselves. Lastly, the law facilitates a partial definition of temperature as a property that, under constant pressure, varies linearly with gas volume. Thus, the law offers a means to explain a particular model of gases to students, identify the features of gases we ought to look out for, and measure the way actual gases behave (as deviations from the ideal), as well as partially explaining other important technical concepts.

---

[5] The *model* account of explanation itself warns against supposing the equivalence of model explanations across fields, recognising that different models draw the counterfactual dependencies in varying ways: some using causal dependences, others structural dependencies, etc. (Bokulich 2011). If anything, it's best not to see model explanations as a single species of explanation, but rather as a family, with many varieties.

According to P&W (2018: 162), these insights about the explanatory functions of chemical laws are equally applicable to logical laws. To show this, they take as a case study Seth Yalcin's (2012) discussion of *modus tollens* (MT), in which a putative counterexample to MT containing probability terms is proposed. In response, Yalcin presents two alternative modal logics to classical logic, both containing probability operators and informational-semantics, which putatively show why (unlike classical logic) such troublesome cases of MT are invalid.

In particular, P&W propose that Yalcin's discussion of MT demonstrates logical laws share two important functional features with the ideal gas law:

Firstly, while the success of Yalcin's new informational-semantic logics are judged by their ability to show why MT fails in the kind of cases involving probability terms he considers, the counterexample "does not serve to put *modus tollens* or the logics containing it into disrepute, as someone who takes truth as an important feature of explanations might expect." Rather, "*modus tollens* and classical logic more generally…are serving a similar explanatory role to the ideal gas law" (P&W 2018: 163).

Secondly, Yalcin's diagnosis of the counterexample shows that logical explanations will "be of the same kind as is found in classical logic," appealing to the same kinds of "factors", using the "right *kinds* of machinery for the construction of the formal language and the model theoretic semantics." Thus, as with the ideal gas law, classical logic holds a privileged role within logic, acting as "inferential scaffolding for Yalcin's [own] account" (P&W 2018: 163).

Thus, according to P&W: (i) MT and other logical rules contained within classical logic serve a similar explanatory role to the ideal gas law (as specified by Woody) and, further, (ii) classical logic holds a privileged role in logic, serving as inferential scaffolding. Yet, contrary to what P&W propose, we have very good reasons to think these supposed similarities are mistaken, given the extent to which they distort logical practice. We'll consider each in turn.

6.1 Similar Explanatory Role

To propose that MT, and other logical rules, play a similar explanatory role to that assigned to the ideal gas law means that: (i) MT is not a descriptive claim about which arguments are valid, but rather a pedagogical prompt that highlights important features of arguments; and thus, contrary to appearances, (ii) Yalcin's putative counterexample is not really a counterexample at all, given that MT properly understood is not a descriptive claim. Two significant challenges face this interpretation of logical rules, given logical practice.

Firstly, if rules of implication were merely useful pedagogical generalisations, rather than descriptions of validity, we would not be able to make sense of attempts to defend a rule from putative counterexamples by *explaining away the counterexample*. Such attempts can take several forms (Martin & Hjortland 2021), including: (i) arguments that, contrary to appearances, the putative counterexample is *not actually* an instance of the target rule (see Lowe's (1987) reply to McGee (1985)); and, (ii) arguments that our initial judgement regarding the putative counterexample are unreliable, due to the complexity of the case or some other confounding variable (e.g., Bledin 2015). Such replies to putative counterexamples are commonplace in the literature. Yet, if the supposed counterexamples constituted no serious challenge to the target rules, as P&W are suggesting, these attempts would be wholly misplaced. Consequently, the mere existence of such attempts to "rescue" these rules from counterexamples suggest that such rules are not *mere* pedagogical prompts for

how practitioners ought to think about arguments. If they were, uncommon exceptions to the rule would be totally expected and accepted.

Secondly, there are research programmes within logic that we cannot make appropriate sense of unless we interpret them as proposing that particular logical rules are *invalid*, and so ought to be rejected. For instance, non-trivialist dialetheists (Priest 2006), who propose that some (but not all) contradictions are true, are *required* to admit that the classically valid rule of explosion is invalid. To do otherwise would commit the dialetheist to trivialism. The same could be said of relevant logicians (Anderson & Belnap 1975), for whom the *rejection* of explosion is required to ensure our correct theory of consequence meets the necessary standards of relevance. Again, these arguments by non-classical logicians treat (classical) rules of implication not as mere pedagogical prompts, but as descriptive claims about what follows from what (i.e., validity).

Thus, unlike the ideal gas law, which according to Woody is maintained by the community *in spite of* the recognition of its failures, recognised failures with a logical rule are often treated as good enough motivation to reject the rule, assuming a workable rival logic not containing the rule exists; hence the occurrence of *rival* non-classical logics. Logical rules appear not to play the same non-descriptive role that Woody assigns to the ideal gas law.

6.2 Classical Logic as Inferential Scaffolding

P&W (2018: 163-4) further propose that classical logic holds a privileged status within logic, "despite [its] known inaccuracy, because of the role [it plays] in establishing standards for intelligibility for logic," just as the ideal gas law does in chemistry. Rather than serving as merely descriptive theories, both provide "inferential scaffolding" for practitioners in their respective fields. In the case of logic, this means that classical logic serves to both: (i) highlight the important features of arguments which logicians ought to pay attention to when evaluating arguments, and (ii) supply logicians with the necessary syntactic and semantic tools to engage in the practice. P&W (2018: 163) interpret the fact that Yalcin uses just the same syntax and model-theoretic "machinery" in his own informational-semantic logics as evidence for these claims.

While correct that contemporary logicians' accounts of validity conform in many respects to that of classical logic, with their use of first-order languages and model theory, it would be a mistake to conclude that this equates to classical logic playing a similar "scaffolding" role to the ideal gas law in chemistry. As noted above, according to Woody's interpretation, the ideal gas law highlights important properties of gases that practitioners ought to pay attention to *irrespective of the law's truth*. So, the perceived truth or falsity of the law is *irrelevant to its privileged position* within practice. The same does not hold true of classical logic. After all, classical logicians spend time still providing arguments *defending classical logic*, whether this to show how their theory can accommodate apparent troublesome cases, such as vague predicates (Williamson 1994), or that it can successfully deflect challenges from competitors, such as intuitionistic logic (Rumfitt 2015). It is difficult to make sense of these activities if classical logic's perceived privileged status were detached from its perceived truth.

Despite this, one may still think that the continued use of first-order languages and model-theoretical semantics are enough to demonstrate that classical logic plays a similar "scaffolding" role to the ideal gas law, specifying the norms that logical explanations ought to adhere to. Yet, this would be a mistake, for two reasons.

Firstly, while contemporary non-classical logics retain certain features of classical logic, others are rejected. To note a few examples, relevant logicians reject the account of validity as truth-preservation, glutty logicians reject the exclusivity of truth and falsity, and others reject model-theoretic semantics in favour of game-theoretic semantics to better model implications from imperfect information (Hintikka & Sandu 1997). Thus, not all features of classical logic have been preserved, providing the framework in which other theories of validity are given.[6]

Secondly, many of the prevalent features of modern theories of validity, though found in classical logic, did *not originate with* classical logic. Rather, they were fruitful features of other theories, or tools, either built upon through the construction of classical logic, or later incorporated into classical logic. For example, the underlying assumption of formal logic that arguments can be classified as (in)valid *in virtue of their form* dates back to Aristotle with the syllogistic tradition, and even the language of propositional logic is found in Stoic works, although the semantics given to the connectives are non-classical (Bobzien 1999). Further, both natural and sequent-calculus proof theory were developed with the analysis of mathematical reasoning within proofs in mind, without a presumption in favour of classical logic (Prawitz 1965).

The picture painted by these cases is that classical logic is best viewed not as a privileged theory, which provides scaffolding for other theories *regardless of its truth*, but as a particularly successful theory which some members of the community think can be improved upon (Martin & Hjortland 2021). If one of the questions that the logical community are interested in, as AEL proposes, is which logic best explains validity, it would come as no surprise if some fruitful features of previous theories were persisted with, and other features deemed to be unsuccessful dropped. Thus, unlike the role that Woody assigns the ideal gas law in chemistry, where its truth is not under question, the preservation of elements of classical logic within competing theories of validity tells not for its privileged status, but rather for its recognised past success in certain domains.

While much more could undoubtedly be said about the differences between explanations in the natural sciences and logic, these brief remarks concerning P&W's analysis serve to demonstrate the importance of resisting the temptation of drawing too close a connection between explanations in the sciences and logic. Just because extra-systemic explanations can be found across various fields does not suffice to conclude that the form these explanations take are indistinguishable. Unless logic's own practices require it, substantiating **VE** does not oblige the anti-exceptionalist to endorse the stronger principle **EE**.

7. Conclusion

This paper set itself the task of providing some clarity and direction to the debate over whether logics explain. It's achieved this by clarifying the type of explanatory practice which is under question, in the form of **VE**, and shown how the truth of **VE** is best tested using the *practice-based approach*, by arguing that the approach falls foul of none of the pitfalls of previous discussions on logical explanations. We have also highlighted, if

[6] Indeed, while many logic textbooks treat classical logic as *the* standard logic in which the mechanisms of logic are presented, to then be deviated from with non-classical logic, this is not always the case. Jan von Plato's (2014) textbook, for example, first introduces the formal techniques of logic using intuitionistic logic, and then describes classical logic as a limiting case within decidable situations.

briefly, an encouraging line of enquiry, drawing a connection between model explanations in other fields and features of logical theories. Our next task is to suitably test **VE** and explore this line of enquiry by looking at logicians' practice in detail.[7]

References

Anderson AR & ND Belnap (1975) *Entailment: The Logic of Relevance and Necessity (Vol. I)*. Princeton, NJ: Princeton University Press.

Ayer AJ (1936) *Language, Truth and Logic*. New York, NY: Dover.

Bledin J (2015) Modus Ponens Defended. *The Journal of Philosophy*, 112: 57–83.

Bobzien S (1999) Stoic Logic. In: K Algra, J Barnes, J Mansfeld and M Schofield (eds.), *The Cambridge History of Hellenistic Philosophy* (pp. 92–157). Cambridge: CUP.

Bokulich A (2011) How Scientific Models Can Explain. *Synthese* 180: 33–45.

Bokulich A (2012) Distinguishing Explanatory from Nonexplanatory Fictions. *Philosophy of Science*, 79: 725–37.

BonJour L (1998) *In Defense of Pure Reason*. Cambridge: Cambridge University Press.

Braillard PA & C Malaterre (2015) Explanation in Biology: An introduction. In: PA Braillard & C Malaterre (eds.) *Explanation in Biology* (pp. 1–28). Dordrecht: Springer

Dalrymple M (2001) *Lexical Functional Grammar* (Syntax and Semantics Series, vol. 34). New York: Academic Press.

De Regt HW (2017) *Understanding Scientific Understanding*. Oxford: Oxford University Press.

Dworkin R (1986) *Law's Empire*. Cambridge, MA: Harvard University Press.

Fagan MB (2015) Collaborative Explanation and Biological Mechanisms. *Studies in History and Philosophy of Science*, 52: 67–78.

Gaukroger S (1989) *Cartesian Logic: An essay on Descartes's conception of inference*. Oxford: OUP.

[7] **Acknowledgements**: I am grateful to colleagues at the University of Bergen for their comments on a draft of this paper, particularly Sorin Bangu, Evelyn Erickson, and Ole Hjortland,. I would also like to thank two anonymous referees for their detailed comments on a previous version of this paper. Research for this paper was supported by a Marie Skłodowska-Curie grant (agreement no. 797507), under the European Union's Horizon 2020 research and innovation programme.

Godden D (2017) Mill on Logic. In: C Macleod & DE Miller (eds.), *A Companion to Mill* (pp. 175–91). Hoboken, NJ: John Wiley & Sons.

Hansson SO (2017, Winter) Logic of Belief Revision. In: EN Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Online: https://plato.stanford.edu/archives/win2017/entries/logic-belief-revision/

Hintikka J & G Sandu (1997) Game-Theoretical Semantics. In: J van Benthem and A ter Meulen (eds.), *Handbook of Logic and Language* (pp. 361–410). Amsterdam: Elsevier.

Hjortland OT (2017) Anti-Exceptionalism about Logic. *Philosophical Studies*, 174: 631–58.

Horn LR (1989) *A Natural History of Negation*. Chicago, IL: University of Chicago Press.

Lowe EJ (1987) Not a Counterexample to Modus Ponens. *Analysis*, 47: 44–7.

Martin B (2021) Identifying Logical Evidence. *Synthese*, 198: 9069–95.

Martin B (forthcoming) The Philosophy of Logical Practice. *Metaphilosophy*.

Martin B & OT Hjortland (2021) Logical Predictivism. *Journal of Philosophical Logic*, 50: 285–318.

Martin B & OT Hjortland (forthcoming) Anti-Exceptionalism about Logic as Tradition Rejection. *Synthese*.

Martini C (2014) Experts in Science: A view from the trenches. *Synthese*, 191: 3–15.

McGee V (1985) A Counterexample to Modus Ponens. *Journal of Philosophy*, 82: 462–71.

Payette G & N Wyatt (2018) How Do Logics Explain? *Australasian Journal of Philosophy*, 96: 157–67.

Prawitz D (1965) *Natural Deduction: A Proof-Theoretical Study*. Stockholm: Almqvist & Wiksell.

Priest G (2006) *In Contradiction: A Study of the Transconsistent*, 2nd ed. Oxford: Clarendon Press.

Priest G (2014) Revising logic. In: P Rush (ed,), *The Metaphysics of Logic* (pp. 211–23). Cambridge: CUP.

Priest G (2016) Logical Disputes and the a Priori. *Logique et Analyse*, 59: 347–66.

Rumfitt I (2015) *The Boundary Stones of Thought*. Oxford: OUP.

Ruphy S (2017) *Scientific Pluralism Reconsidered: A new approach to the (dis)unity of science*. Pittsburgh, PA: University of Pittsburgh Press.

Rossberg M & S Shapiro (2021) Logic and Science: Science and logic. *Synthese*. Online first: https://doi.org/10.1007/s11229-021-03076-w

Resnik N (1999) Against Logical Realism. *History and Philosophy of Logic*, 20: 181–94.

Salmon WC (1971) Statistical Explanation. In: WC Salmon (ed.), *Statistical Explanation and Statistical Relevance* (pp. 29–88). Pittsburgh, PA: University of Pittsburgh Press.

Shapiro S (2000) The Status of Logic. In: C Peacocke and PA Boghossian (eds.), *New Essays on the A Priori* (pp. 333–66). Oxford: Oxford University Press.

Sher G (2016) *Epistemic Friction: An Essay on Knowledge, Truth and Logic*. Oxford: OUP.

Szabó ZG (2012) Against Logical Form. In: G Preyer (ed.), *Donald Davidson on Truth, Meaning, and the Mental* (pp. 105–26). Oxford: Oxford University Press.

Tal E (2020, Fall) Measurement in Science. In: EN Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Online: https://plato.stanford.edu/archives/fall2020/entries/measurement-science/

Tuckwell HC & RJ Williams (2006) Some Properties of a Simple Stochastic Epidemic Model of SIR Type. *Mathematical Biosciences*, 208: 76–97.

von Plato J (2014) *Elements of Logical Reasoning*. Cambridge: Cambridge University Press.

Williamson T (1994) *Vagueness*. London: Routledge.

Williamson T (2017) Semantic Paradoxes and Abductive Methodology. In: B Armour-Garb (ed.), *The Relevance of the Liar* (pp. 325–46). Oxford: Oxford University Press.

Woods J (2019) Against Reflective Equilibrium for Logical Theorizing. *Australasian Journal of Logic*, 16: 319–41.

Woodward J (2003) *Making Things Happen: A Theory of Causal Explanation*. Oxford: OUP.

Woodward J (2017, Fall) Scientific Explanation. In: EN Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Online: https://plato.stanford.edu/archives/fall2017/entries/scientific-explanation/

Woody AI (2015) Re-orienting Discussions of Scientific Explanation: A Functional Perspective. *Studies in History and Philosophy of Science*, 52: 79–87.

Wright C (1986). Inventing Logical Necessity. In: J Butterfield (ed.), *Language, Mind, and Logic* (pp. 187–209). Cambridge: Cambridge University Press.

Wright C (2018) Logical Non-Cognitivism. *Philosophical Issues*, 28: 425–50.

Yalcin S (2012) A Counterexample to Modus Tollens. *Journal of Philosophical Logic*, 41: 1001–24.