# DISSERTATION / DOCTORAL THESIS

Titel der Disseratation / Title of the Doctoral Thesis

„Essays on the Foundations of
Axiom and Logic Selection“

verfasst von / submitted by

Gareth R Pearce B.A. MA

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Doktor der Philosophie (Dr.phil.)

Wien, 2023 / Vienna, 2023

| | |
|---|---|
| Studienkennzahl lt. Studienblatt / degree programme code as it appears on the student record sheet: | UA A 792 296 |
| Studienrichtung lt. Studienblatt / degree programme as it appears on the student record sheet: | Philosophie |
| Betreut von / Supervisors: | Univ.-Prof. Mag. Mag. Dr. Georg Schiemer<br>Assoz. Prof. Mag. Mag. Dr. Dr. Esther Heinrich-Ramharter, Privatdoz. |

# Acknowledgements

I am indebted to a great many people for their support during the production of this thesis.

Professionally, I would like to thank my supervisors Georg Schiemer and Esther Heinrich-Ramharter. Their input throughout the production of this work has been invaluable and the thesis would have been poorer without them. I would also like to thank the members of the *Forms of Normativity: Transitions and Intersections* (FoNTI) project, as well as the wider doctoral community at the University of Vienna, for both their academic input and personal support. I include within this also the many visiting Fellows of the Institute for the Vienna Circle who have discussed my work with me, and the members of the Vienna Forum for Analytic Philosophy, especially in the years 2019-2021. A list of all those from whose input I have benefited would likely exceed the length of this thesis as it stands. I will confine myself to specifically thanking Felix Danowski, for discussions covering near every part of philosophy, Simon Weisgerber, for discussions on the philosophy of mathematical practice, and Marlene Valek, both for a wide range of philosophical discussions and for her help translating the Kurzfassung.

Lastly, in terms of professional acknowledgments, I would like to thank the organizers, attendees and speakers at three summer schools on the philosophy of mathematics held in the summer of 2022: Models & Methods (Vienna), Philosophy of Mathematics Beyond the Armchair?! (HHU Düsseldorf) and the European Summer School in the Philosophy of Mathematics (Konstanz).

Personally, I would, of course, like to thank my friends and family for their support and encouragement of my decision to move to a new country 1500 miles away for my career and studies. I am fortunate enough that a full list of thanks to friends in England, Austria and beyond who have supported me would again be far too extensive to write. I will specifically mention the support of my two closest friends in England: Zak Cohen and Matthew King, both of whom have known and supported me for many years, through all the stages of my study. Lastly, my closest friend in Vienna Ahkin Keizer, who has supported me immeasurably during difficult times in writing this thesis.

# Abstract: English

*Essays on the Foundations of Axiom and Logic Selection* is a collection of six interconnected works on the philosophical foundations of logic and mathematics. In particular, the essays address the kinds of criteria one might apply when selecting one's axioms or logic, and the kinds of philosophical commitments about logic and mathematics which might lead one to prefer one set of criteria over another.

Essays I-IV comprise the essays on axiom selection. They address the question of which criteria mathematicians should apply when evaluating and deciding between competing axiomatic theories, and the philosophical considerations which speak for one theory of axiom selection over another.

Essays V and VI address logic selection. These essays defend the well known, though frequently maligned, Carnapian account of logical correctness.

Three recurrent themes are present through the essays.

I am frequently motivated by an anti-metaphysical approach to philosophy. At several points, perhaps most notably in essays II, III and VI, I argue against more metaphysically loaded theories of axiom and logic selection. The details of this, in particular what I mean by "anti-metaphysical", are summarized to my satisfaction in the overview.

Instrumentalism is the second recurrent theme. I view the value of mathematics and logic as tools for other forms of inquiry, not as the independent study of a particular subject-matter.

Lastly, I defend various versions of pluralism throughout the thesis. In essays I, III and IV I defend (several versions of) axiomatic pluralism. In essay V I defend logical pluralism.

Through these essays, I hope to have made a meaningful contribution to a series of questions that I take to be central to current work in the philosophies of logic and mathematics.

# Kurzfassung: Deutsch

Die Aufsätze über die Grundlagen der Axiom- und Logik-Auswahl ist eine Sammlung von sechs miteinander zusammenhängenden Beiträgen über die philosophischen Grundlagen der Logik und Mathematik. Die Aufsätze befassen sich insbesondere mit den Kriterien, die man bei der Auswahl der Axiome oder der Logik anwenden kann, sowie mit den philosophischen Überzeugungen in Bezug auf Logik und Mathematik, die dazu führen können, dass man eine Reihe von Kriterien einer anderen vorzieht.

Die Aufsätze I-IV beschäftigen sich mit der Auswahl von Axiomen. Sie befassen sich mit der Frage, welche Kriterien Mathematiker*innen bei der Bewertung und Entscheidung zwischen konkurrierenden axiomatischen Theorien anwenden sollten, und mit den philosophischen Überlegungen, die für eine Theorie der Axiomen-Auswahl gegenüber einer anderen sprechen. Die Aufsätze V und VI befassen sich mit der Auswahl der verwendenden Logiken. Sie verteidigen die bekannte, jedoch häufig kritisierte Carnap'sche Theorie der logischen Richtigkeit.

Drei wiederkehrende Themen ziehen sich durch die Aufsätze.

Ich werde häufig von einer anti-metaphysischen Herangehensweise an die Philosophie motiviert. An mehreren Stellen, vor allem in den Aufsätzen II, III und VI, argumentiere ich gegen metaphysisch aufgeladene Theorien der Axiom- und Logik-Auswahl. Die Einzelheiten hierzu, insbesondere was ich unter "anti-metaphysisch" verstehe, sind in der Übersicht für mich zufriedenstellend zusammengefasst.

Der Instrumentalismus ist das zweite wiederkehrende Thema. Der Wert von Mathematik und Logik liegt meines Erachtens in ihrer Funktion als Werkzeuge für andere Formen der Forschung, nicht in unabhängig Erforschung ihres eigenen Wahrheitsbestands.

Abschließend verteidige ich in der gesamten Dissertation verschiedene Versionen des Pluralismus. In den Aufsätzen I, III und IV verteidige ich (mehrere Versionen) des axiomatischen Pluralismus. In Aufsatz V verteidige ich den logischen Pluralismus.

Durch diese Aufsätze hoffe ich, einen bedeutsamen Beitrag zu einer Reihe von Fragen geleistet zu haben, die ich als zentral für die aktuelle Arbeit in der Philosophie der Logik und der Mathematik betrachte.

(Mit Dank an DeepL und Marlene Valek für ihre Hilfe bei dieser Übersetzung)

# Table of Contents

# Publications & Submission

All essays within this collection are currently under review with one of a number of reputable journals.

Essay I: *On Axiom Selection* - Currently under review with *Analysis & Metaphysics*

Essay II: *Mathematics Needs no (Philosophical) Foundation* - Currently under review with *Synthese*

Essay III: *Instrumentalism & Axiom Selection* - Currently under review with the *British Journal for the Philosophy of Science*

Essay IV: *Proofs, Derivations and Axiomatic Pluralism* - Currently under review with *Annals of Mathematics and Philosophy*

Essay V: *Language, Truths and Logics* - Currently under review with the *Philosophical Quarterly*

Essay VI: *Three Approaches to Logical Correctness* - Currently under review with *Logic & Logical Philosophy*

# Overview of Essays on the Foundations of Axiom and Logic Selection

## Introduction

This thesis is a compilation of six essays: four on Philosophy of Mathematics and two on the Philosophy of Logic. Whilst interconnected and presented together, the Essays are intended as stand-alone works and should be evaluated as such.

This overview gives me the opportunity to explain the research context connecting these essays and to draw out and comment on some of the themes present across the thesis as a whole.

I recommend, then, reading this overview twice. Once at the beginning to provide context for the thesis. But then again at the end to summarize the project as a whole.

## Essays I-IV: Essays on Axiom Selection

The four essays on the Philosophy of Mathematics address the topic of axiom selection. Mathematics needs axioms. Without assumptions, the only statements one can prove are simple logical tautologies (e.g. $\vdash \phi \vee \neg\phi$). At least for first-order logics, this is insufficient. In order to do anything that looks like modern mathematics, one needs more substantial assumptions: The Axioms.

But in many areas of mathematics (in particular set theory), one faces a choice between different axiomatizations. Should one, for instance, adopt the axiom of choice?

This leads to a number of questions central to the philosophy of mathematics:

1. What axioms should one use?
2. What criteria should be applied when deciding between axioms?
3. How are the answers to 1. and 2. affected by wider philosophical claims about mathematics?

I do not aim to answer question 1 in this thesis. There has been a great deal of literature dedicated to this question. For instance, the debate between category theory and set theory as a foundation for mathematics[1], or debates about additional axioms in set theory[2].

In my opinion, however, ongoing debates about which axioms to use frequently lack clear criteria for axiom selection. And when such criteria were present, they were frequently not uncontroversial and the kinds of philosophical claims underpinning such views were not always made explicit[3].

The belief behind this thesis is that progress on the first-order question of what axioms to use will be helped by a more theoretical discussion of what the correct theory of axiom selection is. In other words, before establishing what the correct axioms are, one must establish what must be true of some axiom or axiomatic theory for it to be correct; what are the criteria to be applied when selecting one's axioms?

Essay I: *On Axiom Selection* takes one further step of abstraction. Before arguing for my preferred account of axiom selection, it was essential to understand what the possible theories of axiom selection are. The purpose of Essay I is conceptual geography. I hope to have mapped the possibility space of the kinds of theories of axiom selection available, along with some limited comments on the kinds of philosophical claims about mathematics that might lead one to adopting one theory or another.

Essay II: *Mathematics Needs No (Philosophical) Foundation* and Essay III: *Instrumentalism & Axiom Selection* build on Essay I by arguing for and against particular accounts of axiom selection. Essay II presents the negative case against "foundationalist" accounts of axiom selection. Essay III presents my positive account, greatly influenced by the work of Penelope Maddy.

Essay III, in particular, I take to be perhaps the most crucial of the thesis. Two of the collection's wider themes, pluralism and instrumentalism, first seriously appear within this essay (though Axiomatic Pluralism is also discussed in Essay I).

Essay IV: *Proofs, Derivations and Axiomatic Pluralism* stands a little further away from essays I-III, which are all quite interrelated. Essay IV deals with a current and ongoing debate in the Philosophy of Mathematics: the relationship between proofs, in the mathematician's not the logician's sense, and formal derivations. Given the looser relationship between proofs and derivations that has been argued for in recent years[4], I argue for a similarly loose relationship between proofs and axiomatic theories.

This is not per se an essay on axiom selection. Rather, this is an essay on what role axioms actually play in mathematical practice. Philosophers of Mathematics tend to be quite conservative. They typically dislike philosophical positions that are revisionary of the way mathematics is done, and reasonably so. Given this conservatism, a great deal hangs on what role axioms play within mathematical practice. Essay IV, then, is indirectly relevant to the general question of the correct theory of axiom selection.

Essay IV also further develops the theme of pluralism, though in a manner quite different to the pluralism advocated for in essay III. Essay III's pluralism is normative - I argue that there are many axiomatic theories that mathematicians should (or may) use. Essay IV's pluralism is descriptive - I argue that, as a matter of fact, proofs can be associated with multiple derivations from different axiomatic theories.

The aims of the four Essays are therefore as follows:

1. To clarify and taxonomize the possible theories of axiom selection available.
2. To explore some of the philosophical motivations behind differing theories of axiom selection.
3. To make the case for an instrumentalist theory of axiom selection.
4. To make a preliminary case for axiomatic pluralism, the view that there are many correct axiomatic theories.

## Essays V-VI: Essays on Logic Selection

My work on logic selection, logical correctness and logical pluralism began in connection to my work on axiom selection, though the two gradually grew apart.

First, there is some crossover between logic selection and axiom selection with respect to the use of non-classical logics in mathematics. Historically, the primary example of this has been the role of intuitionistic logic within mathematics[5]. More recently, there has been work on applying paraconsistent logics to mathematics[6]. Earlier in my research, I suspected that I may need a combined account of axiom and logic selection. This, I believe, turned out not to be the case.

Second, work on logical correctness and logic selection is far more extensive and developed than its mathematical counterpart[7]. It was my hope that I could build on insights from the philosophy of logic when developing my account of the philosophy of axiom selection.

This bore out to a certain extent. The themes of pluralism and instrumentalism, which became central in my thinking about axiom selection, began squarely within my work on the philosophy of logic. However, there is not, as I'd hoped to find, any simple mapping between theories of logical and axiomatic correctness.

Early in my research into logic selection I became convinced by something adjacent to the Carnapian account of logical correctness[8], if not the Carnapian account itself. The nature of truth and consequently valid inference is language-dependent. Any student of intermediate logic classes learns how to construct languages with non-classical properties. It's a very simple matter, once one knows how, to define a language where, say, disjunctive syllogism is invalid. Picking one's logic is then simply a matter of picking one's language.

My view is not exactly the same as Carnap's. Most notably, I do not accept Carnap's tolerance principle. I refer to the position I defend, Carnap's account of correctness without tolerance, as the Neo-Carnapian view. As an exegetical aside, I happen to think that the importance of the tolerance principle is readily over-stated in interpreting Carnap's work. Carnap appears to violate this principle in multiple places in his wider work, most notably in *Empiricism, Semantics and Ontology* (Carnap, 1950). It might therefore be that my Neo-Carnapian view just is Carnap's real view.

In future work, I hope to argue that this is the case and that we should downplay the importance of the tolerance principle when reading Carnap. But such exegetical questions are put aside in this thesis. This collection is systematic, not historical or exegetical in nature.

Despite what I take to be its natural plausibility, Carnap's account has suffered a great deal of criticism over the last eighty or so years. Much of this criticism is from Quine. Many have even taken the Carnapian position to be all but refuted by Quine[9]. But reports of this position's death have, in my opinion, been greatly exaggerated.

In Essay V: *Language, Truths & Logics*, I defend my Neo-Carnapian position from what I take to be the most common objections levied against it. I hope to have shown that the position is in far better shape than is frequently reported; it therefore is deserving of greater attention within present debates on logical correctness and logical pluralism.

The name of this paper is also a small homage to A.J. Ayer (1936)'s book of a similar name: *Language, Truth and Logic*. His book was one of the first philosophy books I read, having picked up a second-hand copy, aged 17, for about £1 from the local Oxfam bookshop. A small nostalgic nod, I hope, is forgiven.

Essay VI: *Three Approaches to Logical Correctness* provides a very general case for the (Neo) Carnapian view of logical correctness. At a very broad level, I consider three ways one might think about logical correctness: logical realism, what I call the one-language view and the Carnapian tradition. I present the Carnapian objections to the former two positions.

The generality of this essay comes out of two competing concerns. On the one hand, I took it to be important that I provide a positive case for my view of logic, not merely defend it. On the other hand, however, there are simply too many views of logic out there to do justice to all the available positions in a short space. This thesis is, primarily, a work on the philosophy of mathematics, not the philosophy of logic. It would have been too great a tangent from the intended purpose of the work to respond in detail to the vast span of positions presently defended.

Essay VI, then, is a compromise. I don't claim to have definitively argued for the (Neo) Carnapian position. I hope, however, to have given a sense, a flavour, of the kinds of arguments the (Neo) Carnapian would make in response to certain kinds of positions.

Returning to the context of the thesis as a whole, adopting a Carnapian view of logic had the ironic consequence of severing any strong connection between the two parts of the thesis. A logical realist is, perhaps, compelled to use whichever logic they take to be objectively correct within mathematics. My Neo-Carnapian is under no such obligation. For me, selecting the correct logic for mathematics is simply a pragmatic question of selecting the most functional language for mathematics. But this is then simply congruent with my account of axiom selection in general.

In summary, then, the aims of Essay V and Essay VI are as follows:

1. To outline and defend logical pluralism on the basis of a (Neo) Carnapian account of logical correctness.
2. To defend the Carnapian account from nearly a century of sustained criticism.
3. To outline some of the benefits of the Carnapian view over some of its rivals.

# Philosophical Themes across the Essays

Though the six essays of this thesis are intended to stand alone, there are several themes which run through the thesis as a whole, connecting the the works together. In this section, I provide some brief comments on each.

## Anti-Metaphysical Motivations

In my MA Thesis, I defended epistemic nominalism about Mathematics. I argued that even if there are objective mathematical truths, we could not have any knowledge of them. I have also publicly defended nominalism by offering a nominalist reply to Linnebo (2018)'s argument for abstract objects in *Thin Objects* (Pearce, 2022).

Earlier in my research, I had similar concerns regarding logical realism. These concerns rested, to a certain extent, on an overly narrow conception of logical realism.

There are certainly versions of logical realism that do fall fowl of epistemic nominalist worries but there are also versions of the view that do not. Maddy (2007)'s view of logic, for instance, is realist but nominalist. I discuss the relationship between nominalism and logical realism briefly in Essay VI, though leave a great deal unsaid.

Nevertheless, during my research, my nominalism grew into a more general weariness of overly metaphysical philosophy. In particular, even if there are objective metaphysical facts and we can have knowledge of them (both of which I'm sceptical), I was unclear about the relevance of these facts to the actions of logicians and reasoners.

Parallel arguments appear in Essays II and VI relating to this.

In Essay II I ask, even if there are mathematical facts, why should mathematicians care? If we discovered, tomorrow, that actual continua are structured in just such a way that makes CH true, why would that mean that set theorists should adopt CH?

In Essay VI I ask, even if there are logical facts, why should we care for the purposes of our practical or scientific reasoning? Even if, fundamentally, the world has Boolean structure, why should that mean that we dispense with fuzzy predicates for our practical purposes?

These arguments are instances of my anti-metaphysical approach to theoretical philosophy generally, which owes some degree of debt to the (ongoing) program of conceptual engineering. When it comes to understanding most of the objects or concepts of theoretical philosophy (knowledge, justification, explanation, grounding, causation, etc), I am less concerned with finding objective answers to these questions than I am with understanding what practical work these concepts need to do as part of our inquiry, and adjusting our concepts to best fit the needs of that work.

That being said, I did not wish to premise my thesis on a nominalist or anti-metaphysical worldview. These wider ideological commitments are a product of the arguments present in this collection, not the groundwork for them.

Lastly, even if one disagrees with these commitments there is still, I hope, enough of interest in the thesis. Essay I and Essay IV, in particular, stand apart from these commitments.

## Instrumentalism about Logic and Mathematics

In the summer of 2022 I had a realization that changed, I think quite significantly, how I think about axiom selection. I realized that a claim I had been making for quite some time was, if not entirely wrong, at least subtly missing the mark.

During the early and middle stages of my research, I believed that one's views on the nature of mathematical truth and knowledge would all but determine one's view of axiom selection. I hoped to show, more or less, that Fregian-style realists were committed to an epistemic theory of axiom selection (*"Epistemic Foundationalism"* in Essay II's terminology), Quine-style realists were committed to a descriptive theory of axiom selection (*"Ontic Foundationalism"*) and nominalists were committed to a normative theory of axiom selection (like my own outlined in essay III, influenced heavily by Maddy (2011)).

This is not entirely wrong and most realists of those types and nominalists will accept something like those positions. What I struggled to explain were realists who adopted normative theories of axiom selection. Maddy, in particular, was difficult to categorize. She is a realist but has a normative theory of axiom selection.

I had the fortune that summer of attending three summer schools on the philosophy of mathematics in Vienna, Düsseldorf and Konstanz in turn. This gave me extensive opportunities to discuss my work with some exceptionally talented individuals. In these conversations, I found myself more and more explaining my views with reference not to my views on mathematical truth, but rather my views on the *purpose* of mathematics.

This I now take to be perhaps the most major fault line in dividing theories of axiom selection. Whilst Maddy is a realist, her realism is somewhat coincidental. Finding true mathematics is secondary to finding *useful* mathematics. Truth, for both Maddy and me, plays no role in understanding the value of mathematics. For realists unlike Maddy, however, the purpose of mathematics is the discovery of the objective mathematical truths.

I suspect there are more views about the function of mathematics than the instrumentalist vs descriptivist dichotomy that I outline. Maddy's instrumentalism and my own, for instance, differ quite substantially. But I hope to have identified, at least, the two poles in this debate.

Of course, there are going to be overlaps in one's view of mathematical truth and one's view of the purpose of mathematics. My earlier view missed the mark but was not entirely unreasonable. But it's one's view of the purpose, function or value of mathematics that I now take to be most consequential in determining the correct theory of axiom selection.

This is interesting for two reasons.

First, it shows that debates over theories of axiom selection don't simply boil down to traditional, well-trodden battle lines between nominalists and realists. There's a different sort of disagreement at play here.

Second, whilst there has been more written on mathematical truth than one could read in a lifetime, there has been sparsely anything written, at least within anglophone analytic philosophy, on the value, function or purpose of mathematics. To my knowledge, there is no book-length discussion of this question.

This normative question about the purpose of mathematics I take to be the most interesting point for future research arising from these essays.

Lastly, to connect this theme to the essays on logic selection, the Neo-Carnapian view of logic selection is quite clearly instrumentalist in nature. Logical truth comes easily - it's just a matter of picking one's language appropriately. The interesting question is what's most useful for our logical concepts or terms to do. But the value of logic, then, is instrumental. It's not about having a descriptively correct theory of the logical facts, but rather having a useful logic that plays the role we'd like it to.

My instrumentalism about logic and mathematics, then, is closely tied to my anti-metaphysical motivations outlined above.

## Pluralism in Logic and Mathematics

The final theme I wish to explain in this overview is pluralism. My interest in finding the best account of logical correctness came out of my interest in logical pluralism. In the early stages of this project, one of my research questions was if the insights behind logical pluralism could be applied to theories of axiom selection to motivate axiomatic pluralism.

This particular line of inquiry did not bear fruit to quite the extent that I'd hoped. There was no simple story to tell where one could simply walk one's favourite theory of logical pluralism into the mathematical domain and, without much adaptation, have an account of axiomatic pluralism.

Nevertheless, both my axiomatic pluralism, as defended in essays I, III and IV and my logical pluralism, as defended in Essay V, do share common motivations. Both, I think, are natural consequences of the Instrumentalism outlined above. The descriptivist must concern themselves with finding the "One True Logic" or "One True Axiomatic Theory". They can only be pluralist if reality itself turns out to be pluralist (a puzzling proposition). But an instrumentalist has no such constraints. Logic and mathematics are tools for certain kinds of representational activities. There might be different representational tools appropriate for different purposes. Consequently, one needs many logics and many axiomatic theories.

My pluralism, then, is a consequence of my instrumentalism.

Within my work, I'm very keen to stress that there are often very many types of pluralism that can go by the same name. One should never claim to be a logical pluralist, alethic pluralist, axiomatic pluralist or any suchlike without qualification. One must always be a monist, pluralist or nihilist relative to a specific number-question that asks how many of some kind of thing there is. Often, subtly different formulations of the question can impact the outcome.

The question "how many logics are correct?" is a prime example of this. The question itself is ambiguous as it doesn't specify the relevant notion of correctness. Different ways of specifying this will give different results. One might have very different answers to the questions "How many logics are objectively correct?", "How many logics are correct for some language?" and "Given some language, how many logics could be correct?"

It's important when claiming to be a monist, pluralist or nihilist to be really quite precise about the question to which one answers one, many or none.

There are two important points of differentiation between similar but distinct number-questions that appear at multiple points in these essays.

The first is between descriptive and normative formulations of questions. "How many logics are correct?" and "How many logics may we reason with?" are not equivalent questions. It's presumptive to assume that truth has pro-toto normative force within a given context. This distinction plays an important role in Essay V in relation to logical pluralism, but also in essay IV in relation to axiomatic pluralism. I do not defend the same notion of axiomatic pluralism across all the essays. The versions of axiomatic pluralism outlined in Essay I, one of which is defended in Essay III, is normative. Essay IV's axiomatic pluralism, however, is descriptive. It asks how many axiomatic theories *actually* serve as foundations for mathematics.

The second point of differentiation is in specifying the domain or context within which one evaluates. Carnap, for instance, thinks that there are no logics true of all languages, many true of some, and precisely one true of any given language. How one specifies the scope of the number-question turns Carnap into a nihilist, pluralist or monist respectively. Similarly, it's important to differentiate pluralism about the general foundations of mathematics from pluralism about specific mathematical domains. One might, I think reasonably, claim that individual mathematical domains need singular axiomatic theories. Perhaps the role of these theories is precisely to *define* the domain. But nevertheless one might reject the need for a singular foundational theory in which to do meta-mathematics.

Lastly, on pluralism, there is often a tendency to search for the right kind of pluralism. Certain pluralist views are disparaged because they do not get the right kind of pluralism. I certainly don't wish to say that all types of pluralism are equally interesting, they are not, but several objections against Carnapian logical pluralism involve saying that some particular version of pluralism is the *right* notion of pluralism, then showing how the Carnapian is not pluralist in that regard.

I encourage a little more pluralism about pluralism. Whilst different pluralist thesies may be more or less interesting than one another, there's no need to identify the One True Pluralism. One can call oneself a pluralist about logic or axioms without committing to pluralism about every number-question relating to that topic. One might be pluralist in some ways, and monist or nihilist in others. One might liberally call one's self a pluralist if at least some of the number-questions about which one is a pluralist are interesting and worthy of concern, but clarify in more thorough contexts the exact ways in which one thinks pluralism is true.

## Concluding Remarks

In these essays, I hope to have made meaningful progress on a number of questions I take to be central to the Philosophy of Mathematics and Philosophy of Logic. I hope, in particular, to have better clarified the kinds of criteria one might use as part of a theory of axiom selection. I hope also to have clarified the kinds of philosophical claims about mathematics, in particular the function of mathematics, which might influence one's account of axiom selection. Adjacently, I hope to have put some life back into the typically disregarded Carnapian account of logic. I have, I hope, addressed the major concerns regarding this position and provided something of a positive account in its favour. Lastly, across both the essays on logic and the essays on mathematics, I hope to have emboldened the prospects for pluralism and instrumentalism within these two domains.

# Notes

[1]See Univalent Foundations Program (2013) and Corfield (2020) for formal introductions to various versions of homotopy type theory. See Mac Lane (1986), Feferman (1977), Mayberry (1977), Hellman (2003), Awodey (2004), McLarty (2004), Linnebo and Pettigrew (2011) and various works by Ladyman & Presnell including Ladyman and Presnell (2016), Ladyman and Presnell (2019) and Ladyman and Presnell (2020). More recently see several chapters in the edited collection Centrone et al. (2019).

[2]Again, see the Centrone et al. (2019) collected volume. See Maddy (1993) and Maddy (2011), the collected works of: Feferman (2000), Steel (2000), Maddy (2000) and Friedman (2000). See also discussion of the multiverse program in Hamkins (2012) and Maddy and Meadows (2020).

[3]A notable exception to this is Maddy (2011)'s work in *Defending The Axioms*, also earlier in Maddy (2007), *Second Philosophy*. To date, her work is the most developed on axiom selection and has played a central role in the development of my thinking about these questions.

[4]See Azzouni (2004), Burgess and Toffoli (2022) Toffoli (2021), Hamami (Hamami (2014) and Hamami (2022)), Hamami and Morris (ming), and Tanswell (Tanswell (2015) and Tanswell (2016)).

[5]Brouwer (1992) and Brouwer (1981), for the classical view. See Iemhoff (2020) and Bridges et al. (2022) for an overview and discussion.

[6]See Weber (2010a), Weber (2010b), Weber (2012), McKubre-Jordens and Weber (2012), Meadows and Weber (2016) and Badia et al. (2022). Also, see Weber's book *Paradoxes and Inconsistent Mathematics* (Weber, 2021).

[7]See Beall and Restall (2005), Shapiro (2006), Priest (2006), Field (2009) and Griffiths (2022). See Russell (2021) and Cook (2010) for overviews.

[8]Carnap (1928), Carnap (1934), Carnap (1947) and Carnap (1950)

[9]Quine (1936) and Quine (1970)

# References

Awodey, S. (2004). An answer to hellman's question: Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 12(1):54–64.

Ayer, A. J. (1936). *Language, Truth and Logic*. Victor Gollancz Ltd, London.

Azzouni, J. (2004). The derivation-indicator view of mathematical practice. *Philosophia Mathematica*, 12(2):81–106.

Badia, G., Weber, Z., and Girard, P. (2022). Paraconsistent metatheory: New proofs with old tools. *Journal of Philosophical Logic*, 51(4):825–856.

Beall, J. and Restall, G. (2005). *Logical Pluralism*. Oxford, England: Oxford University Press.

Bridges, D., Palmgren, E., and Ishihara, H. (2022). Constructive Mathematics. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2022 edition.

Brouwer, L. (1981). *Brouwer's Cambridge lectures on intuitionism*. Cambridge University Press, Cambridge.

Brouwer, L. (1992). *Intuitionismus*. Wissenschaftsverlag, Mannheim.

Burgess, J. and Toffoli, S. D. (2022). What is mathematical rigor? *Aphex*, 25:1–17.

Carnap, R. (1928). *Der Logische Aufbau der Welt*. Hamburg: Meiner Verlag.

Carnap, R. (1934). *Logische Syntax der Sprache*. Springer, Vienna. Translated into English 1937 by Amethe Smeaton under the title "Logical Syntax of Language", Routledge, London.

Carnap, R. (1947). *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago, IL, USA: University of Chicago Press.

Carnap, R. (1950). Empiricism, semantics and ontology. *Revue Internationale de Philosophie*, 4(11):20–40.

Centrone, S., Kant, D., and Sarikaya, D. (2019). *Reflections on the Foundations of Mathematics: Univalent Foundations, Set Theory and General Thoughts*. Springer Verlag.

Cook, R. T. (2010). Let a thousand flowers bloom: A tour of logical pluralism. *Philosophy Compass*, 5(6):492–504.

Corfield, D. (2020). *Modal Homotopy Type Theory: The Prospect of a New Logic for Philosophy*. Oxford, England: Oxford University Press.

Feferman, S. (1977). Categorical foundations and foundations of category theory. In Butts, R. and Hintikka, J., editors, *Logic, Foundations of Mathematics and Computability Theory*, pages 149–169. Reidel, Dordrecht.

Feferman, S. (2000). Why the programs for new axioms need to be questioned. In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Field, H. (2009). Pluralism in logic. *Review of Symbolic Logic*, 2(2):342–359.

Friedman, H. M. (2000). Ordinary mathematics will need new axioms. In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Griffiths, O. (2022). *One True Logic: A Monist Manifesto*. Oxford University Press.

Hamami, Y. (2014). Mathematical rigor, proof gap and the validity of mathematical inference. *Philosophia Scientiae*, 18(1):7–26.

Hamami, Y. (2022). Mathematical rigor and proof. *Review of Symbolic Logic*, 15(2):409–449.

Hamami, Y. and Morris, R. (forthcoming). Rationality in mathematical proofs. *Australasian Journal of Philosophy*.

Hamkins, J. D. (2012). The set-theoretic multiverse. *Review of Symbolic Logic*, 5(3):416–449.

Hellman, G. (2003). Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 11(2):129–157.

Iemhoff, R. (2020). Intuitionism in the Philosophy of Mathematics. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2020 edition.

Ladyman, J. and Presnell, S. (2016). Does homotopy type theory provide a foundation for mathematics? *British Journal for the Philosophy of Science*, page axw006.

Ladyman, J. and Presnell, S. (2019). Universes and univalence in homotopy type theory. *Review of Symbolic Logic*, 12(3):426–455.

Ladyman, J. and Presnell, S. (2020). The hole argument in homotopy type theory. *Foundations of Physics*, 50(4):319–329.

Linnebo, O. (2018). *Thin Objects: An Abstractionist Account*. Oxford: Oxford University Press.

Linnebo, O. and Pettigrew, R. (2011). Category theory as an autonomous foundation. *Philosophia Mathematica*, 19(3):227–254.

Mac Lane, S. (1986). *Mathematics: Form and Function*. Springer-Verlag, New York.

Maddy, P. (1993). Does v. equal l? *Journal of Symbolic Logic*, 58(1):15–41.

Maddy, P. (2000). Does mathematics need new axioms? In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Maddy, P. (2007). *Second Philosophy: A Naturalistic Method*. New York: Oxford University Press.

Maddy, P. (2011). *Defending the Axioms: On the Philosophical Foundations of Set Theory*. Oxford, England: Oxford University Press.

Maddy, P. and Meadows, T. (2020). A reconstruction of steel's multiverse project. *Bulletin of Symbolic Logic*, 26(2):118–169.

Mayberry, J. P. (1977). On the consistency problem for set theory: An essay on the cantorian foundations of classical mathematics (i). *British Journal for the Philosophy of Science*, 28:1–34.

McKubre-Jordens, M. and Weber, Z. (2012). Real analysis in paraconsistent logic. *Journal of Philosophical Logic*, 41(5):901–922.

McLarty, C. (2004). Exploring categorical structuralism. *Philosophia Mathematica*, 12(3):37–53.

Meadows, T. and Weber, Z. (2016). Computation in non-classical foundations? *Philosophers' Imprint*, 16.

Pearce, G. R. (2022). A nominalist alternative to reference by abstraction. *Theoria*, 1:1–12.

Priest, G. (2006). *Doubt Truth to Be a Liar*. New York: Oxford University Press.

Quine, W. V. (1936). Truth by convention. In *Philosophical Essays for Alfred North Whitehead*, pages 90–124. London: Longmans, Green & Co.

Quine, W. V. O. (1970). *Philosophy of Logic*. Cambridge, Mass.: Harvard University Press.

Russell, G. (2021). Logical Pluralism. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2021 edition.

Shapiro, S. (2006). *Vagueness in Context*. New York: Oxford University Press.

Steel, J. R. (2000). Mathematics needs new axioms. In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Tanswell, F. (2015). A problem with the dependence of informal proofs on formal proofs. *Philosophia Mathematica*, 23(3):295–310.

Tanswell, F. S. (2016). Proof, rigour and informality : A virtue account of mathematical knowledge. In *St Andrews Research Repository Philosophy Dissertations*.

Toffoli, S. D. (2021). Groundwork for a fallibilist account of mathematics. *Philosophical Quarterly*, 7(4):823–844.

Univalent Foundations Program, T. (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. `https://homotopytypetheory.org/book`, Institute for Advanced Study.

Weber, Z. (2010a). Extensionality and restriction in naive set theory. *Studia Logica*, 94(1):87–104.

Weber, Z. (2010b). Transfinite numbers in paraconsistent set theory. *Review of Symbolic Logic*, 3(1):71–92.

Weber, Z. (2012). Transfinite cardinals in paraconsistent set theory. *Review of Symbolic Logic*, 5(2):269–293.

Weber, Z. (2021). *Paradoxes and Inconsistent Mathematics*. Cambridge University Press.

# Essay I: On Axiom Selection

An enduring type of debate in the philosophy of mathematics, and in fact within mathematics itself, is over which axioms ought be used. 18th and early 19th-century geometers debated the parallel postulate. The 19th century saw the axiomatisation of group theory. The late 19th and early 20th centuries saw the debate between classical and constructive mathematics. Early set theorists debated the axiom of choice. Debates over the continuum, large cardinals or types of determinacy continue to this day. Some have even suggested the adoption of a set-theoretic multiverse. Outside of set theory, alternative foundations have been proposed, most notably Homotopy Type Theory (Univalent Foundations Program, 2013).

Despite all this, there's far less work outlining or defending *criteria* for axiom selection; the conditions that an axiom system needs to meet in order to be correct. The game of axiom selection is being played without a concrete understanding of the rules. There has been insufficient discussion over the correct theory of axiom selection. The purpose of this paper is to begin to address this gap.

I do not, however, try to do this by arguing for a particular theory of axiom selection. Any case for one theory of axiom selection over another will be highly non-neutral with respect to wider philosophical questions about the metaphysics, epistemology, axiology and purpose of mathematics. What a nominalist like myself with a very instrumental view of mathematics would say about a theory of axiom selection would differ greatly from what a realist would say.

Instead, the aim is to complete what I've taken to calling a *conceptual geography* of theories of axiom selection. Rather than argue for one theory over another, this paper tries to map the logical space between theories; to understand the fault lines and disagreements. Even if philosophers of mathematics are not of agreement regarding the correct theory of axiom selection, my hope is that we can nevertheless clarify precisely where disagreement takes place.

A secondary conclusion is that axiom selection is not a philosophically neutral activity. Any theory of axiom selection is woven with assumptions about the nature and purpose of mathematics. There is a rich *philosophy* of axiom selection that has not yet been properly mapped or understood. However, this is not so much a conclusion that can be argued for, rather it needs to be shown. At several points in the conceptual map, I identify how wider questions in the philosophy of mathematics have an impact on one's theory of axiom selection. Whilst there is no section specifically dedicated to this thesis, it should nevertheless be a claim that's clear and present throughout the paper.

As a final introductory comment, this paper is not intended as a final and complete mapping of the possibility space. In cartography, it is common to first draw a more elementary map, before improving on it, adding details over time. So too with conceptual geography. With time, important distinctions not raised in this paper may become recognised and appreciated. Similarly, some distinctions within this paper might become less relevant and important as our collective understanding of the issues grows. My aim here is to draw the elementary map, to make meaningful and productive progress towards understanding the possibility space, but certainly not to give the final word.

§1 of this paper explains some preliminaries and background regarding axiom selection: What are axioms? For what and in what contexts might axioms be selected? This draws heavily on Dirk Schlimm's work (Schlimm, 2013). §2 maps and taxonomizes possible evaluation criteria for axiom systems. It draws out the distinction between four types of evaluation criteria for axiom systems: ontic or metaphysical, epistemic, conceptual and normative or pragmatic.

With the conceptual geography completed in §2, §4 considers a natural corollary: the question of axiomatic pluralism. Despite the present popularity of pluralist projects across philosophy, pluralism in the philosophy of mathematics remains a comparatively underdeveloped topic. A natural way one might apply pluralism to the philosophy of mathematics is to be a pluralist about the number of correct axiom systems. Building on §2 as well as the literature on logical and alethic pluralism, §3 considers a number of ways in which one might be an axiomatic pluralist. Again, the aim here is not to argue for any particular kind of axiomatic pluralism, but rather to map the space of the kinds of pluralist views available.

In summary:

§1: What are axiom systems, what are they like and when do we pick them?

§2: Evaluation criteria for axiom systems

§3: Axiomatic Pluralism

# 1 What are axiom systems, what are they like and when do we pick them?

## 1.1 What are axiom systems?

The account of axioms one will find in a standard logic textbook treats axioms as a set of propositions. An axiom is an individual mathematical statement. An axiomatic system (or theory) is simply a set of such statements. Often, this is restricted to a set of finitely many axioms and axiom schema[1].

A first possible point of philosophical disagreement is over the admissibility of infinitary axiom systems, i.e. axiom systems that are not stateable via finitely many axioms or axiom schemes[2]. Of course, one can have infinitary theories: $Th(\mathbb{N})$, for instance. But if such theories could take the honorific "axiomatic system" is another question. For instance, Corcoran (1973) argues against the use of infinitary axiom systems.

There are two options Corcoran might take here. First, he might bake into the definition of an axiom system that axiom systems need to be finite. Second, he might allow for the possibility of infinitary axiom systems, but reject them on the evaluative level (see §2). The latter is probably closer to what he actually had in mind.

[1]Axiom schema are an infinite set of procedurally generateable axioms. Typically, one might iterate over all definable formulae with some property, e.g. a certain number of free variables. The induction scheme for arithmetic, for instance, is really an induction axiom for each formula of one free variable

[2]This is not to be confused with Finitism (Zach, 2023), the view relating mathematical truth to the termination of certain finite procedures.

The "sets of statements" definition is sufficient for the purpose of ordinary model theory or mathematical logic. A slightly broader definition is needed for the present philosophical context. Four ways the definition needs to be expanded: notion of consequence, independence of axioms, language and interpretation.

First, the notion of consequence. Schlimm (2013) defines axiomatic systems as sets of statements plus a logical consequence relation. A logical consequence relation is a relation between sentences indicating which inferences are valid or invalid. This will typically be a procedurally generated relation on the basis of a number of basic inference rules, shown to be sound for some language.

The inclusion of a consequence relation in one's notion of an axiomatic system is very understandable, given certain historic debates on axiom selection. The debate between classical and intuitionistic mathematics, for instance, is not about the admissibility of one axiom over another, but rather the admissibility of a particular type of proof (proof by contradiction). To capture disagreements such as these, following Schlimm, logical consequence needs to be included as part of the axiomatic system, and at stake in axiom selection.

There's a complex question about the relationship between logic selection and axiom selection. Does one's logic need to be settled prior to one's selection of axioms? Do mathematical considerations weigh on matters of logic selection? Does one's logic for mathematics need to be one's logic for day-to-day life or for philosophy? Whilst interesting, these questions would require a large diversion into the philosophy of logic and will consequently be avoided here. What matters for present purposes is that settling one's notion of consequence is part of axiom selection. The question of the relative priority of the two is left for another day.

Second, the independence of the axioms. A common at least desirable feature of axiom systems is the mutual independence of the axioms themselves. Of course, this is not a question that's independent of one's choice of a notion of consequence. As with the question of infinitary axiomatic systems, one might choose to include the requirement of independence within the definition of an axiom system, or one might choose to exclude it, but take independence to be a desideratum applied at the evaluative level.

Something also at stake in one's choice of axiomatic system is the language one uses. Corfield (2020), for instance, argues that the language of homotopy type theory (HoTT) has certain benefits over the language of first-order logic. Sometimes these advantages should really be understood via the consequence relation. Adding higher-order variables and quantifiers to the language facilities using a stronger, higher-order logic. The linguistic change facilitates an improvement of the consequence relation. Other times this is not the case, for instance in Corfield's discussion of so-called donkey sentences. In this case, Corfield gives an example of a natural language sentence more easily passed in the language of category theory than first-order logic. Thus the language of one's axiomatic system is also a determinable of the system.

Finally, there's a more philosophically laden possible feature of axiomatic systems: the question of the *interpretation* of the symbols used in stating the axioms.

As the definition stands, it's sufficient for something to be an axiomatic system for it to be an uninterpreted set of statements with a consequence relation. This kind of an axiomatic system would be 'mere' syntax, without any richer meaning. But one might reasonably demand that this is not enough. The axioms cannot just be syntactic, their symbols need an interpretation. For example, there might be some independent, perhaps pre-mathematical, notion of membership to which the '$\in$' symbol refers. Similarly, one might think that there is some real object to which '$\varnothing$' refers.

Again, this brings up a number of philosophical questions:

First, do interpretations of the symbols used in one's axioms need to be provided?

I think there are three natural responses to this question. I dub them the formalist, platonist and structuralist responses, but this is just to give a sense of the kinds of views that might motivate these responses. I don't mean to say that formalism, platonism or structuralism, the metaphysical positions, *entail* these responses; they're responses in the style of the respective views.

The three responses:

Formalist: The symbols need no interpretation, the syntactic information (e.g. arity of the predicates) is enough. This is because mathematics is merely the manipulation of these symbols by logical rules, interpretation plays no role

Platonist: The symbols need an independent interpretation. This is because axioms need to be true and only interpreted sentences can be true.

Structuralist: The symbols need an interpretation, but the interpretation is provided by the axioms themselves. In lattice theory, for instance, “<” is given the interpretation by the axioms of an order relation, but needs no external interpretation beyond that. Perhaps “<” is an arbitrary term for order relations in general. Perhaps some reference-magnetism story can be told and “<” ends up picking out a particular order, it’s just unknown which.

What’s interesting to note here is the direction of reasoning. In both the formalist and platonist cases, claims about how axiomatic systems are to be evaluated is influencing what the conception of an axiomatic system should be. The platonist, for instance, desires their axioms to be true. In virtue of this, their axiomatic systems need to be something truth-apt: interpreted, rather than uninterpreted, symbology.

Once again, though, there’s the option to handle these worries at the level of evaluation, not definition. A platonist might allow a merely syntactic set of statements to count as an axiomatic system (i.e. to adopt the formalist’s definition) but to say that any uninterpreted axiomatic system will fail their evaluation criteria, as truth is a necessary condition for selection and uninterpreted statements aren’t truth-apt.

The second question arising out of the question of interpretation is what the content of a possible interpretation would be. Above, I uncritically switch between symbols being attached to concepts and attached to objects or relations. The difference is quite substantial.

Concepts are cognitive objects. The concept of X does not entail the existence of X. Even a nominalist might grant that the terms of an axiomatic system should be interpreted with mathematical *concepts*, these are just concepts either without or with no known extension.

Objects and relations, however, exist in the world. A platonist might insist that some term $t$ in their axioms actually refers to some real mathematical object $o$. To require this would very much be to bake one's realism into one's concept of an axiomatic system from the ground up.

Thus, this is the first point of substantial possible disagreement about the nature of axiomatic systems. Do systems need to be interpreted? If so, what are the objects of their interpretation? Concepts or objects?

The following, then, are the determinables of an axiom system, with relevant points of philosophical disagreement marked:

1. A set of non-logical statements (the axioms themselves)
   - Do the statements need to be finitely many axioms and axiom schema?
   - Do the axioms need to be independent?
2. A consequence relation
   - What's the relation between axiom and logic selection?
3. A language in which the axioms are stated
4. An interpretation of the axioms
   - Is an interpretation necessary? Can uninterpreted sentences be axioms?
   - With what are symbols interpreted: concepts or objects?

## 1.2 The 'classical' definition of an axiom

The 'classical' definition of an axiom is something like the following:

1. A statement or proposition which is regarded as being established, accepted, or self-evidently true (Google-Oxford Languages)
2. An established rule or principle or a self-evident truth (Merriam-Webster #2)
3. A maxim widely accepted on its intrinsic merit (Merriam-Webster #3)

This 'classical' definition has three parts:

1. Axioms are true. (Metaphysical Claim)
2. Axioms are self-evident. (Epistemic Claim)
3. Axioms are accepted or established. (Group-Doxastic Claim)

It will not be controversial to say that most, if not all, philosophers of mathematics and mathematicians have moved away from this definition[3]. Mayberry (1994) (see also Mayberry (2000)) is the closest in recent years to endorsing something like this definition, but even he backs off of it as an unrealizable ideal.

There is, however, one important difference worth pulling out between this paper's definition and the classical definition.

In the dictionary definitions, "axiom" is a success term. For something to be an axiom it has to succeed by metaphysical, epistemic and doxastic criteria. The evaluative claims that an axiom system is successful in these regards is packaged into the definition of an axiomatic system from the start.

This is obviously unhelpful if the evaluative status of various axioms or axiom systems is precisely what's under discussion. As a result, I avoid using the terms 'axiom', 'axiom system' or any other related conjugates as success terms.

This has bearing on some of the philosophical concerns raised in the previous section. I don't require axiom systems *by definition* to be finitary, independent or interpreted. Now, these might nevertheless be important considerations when *evaluating* axiom systems. But it's unhelpful to pack the evaluation criteria into the definition.

[3]The Oxford Dictionary of Philosophy, for instance, departs wildly from the 'classical' definition: "A proposition laid down as one from which we may begin; an assertion that is taken as fundamental, at least for the purposes of the branch of enquiry in hand. The axiomatic method is that of defining a set of such propositions, and the proof procedures or rules of inference that are permissible, and then deriving the theorems that result." This is, in fact, not too far from this paper's definition.

## 1.3 In which contexts does Axiom Selection take place?

Dirk Schlimm has an extensive paper detailing the role of axiomatisation in mathematical practice. This section draws heavily on §1 and 2 of his paper. §4 of Schlimm is relevant for the next section of this paper.

One feature of axiomatics Schlimm discusses that is not covered in this paper is the *diachronic* role of axiomatisation in the epistemology of mathematics. One of the main points Schlimm makes, which I entirely agree with, is the claim that axiomatization plays a role in mathematical knowledge generation. Schlimm argues against a prior view that axiomatization is merely the systematization of an already well-formed and understood domain (Copi (1958), **?**, Hanson (1958a) and Hanson (1958b)). There's no new knowledge generated at the point of axiomatization. Another way to put this is that axiomatization really belongs in the context of justification, not the context of discovery. Schlimm argues against this prior view. For instance, in axiomatizing geometry, the independence of the parallel postulate was discovered. This led to the discovery of non-Euclidian geometries.

I agree with Schlimm entirely on this matter, but it's a point outside the scope of this paper. This paper is about the possible criteria for deciding between axiomatic systems, not about the relationship between axiomatization and mathematical knowledge.

Within which contexts and for what ends are axioms chosen?

There are three distinctions which are worth drawing out here: pure vs impure, descriptive vs prescriptive and local vs global. With the exception of the descriptive vs prescriptive distinction, these terms are my own and are explained below.

As a point of clarification here, these distinctions are distinctions between *contexts* in which one might find one's self choosing between rival axiom systems, not between types of axiom systems themselves.

The local vs global distinction, for instance, is the difference between selecting an axiom system for some specific mathematical domain (e.g. geometry, group theory or arithmetic) and selecting one as a foundation for the whole of mathematics. This isn't a property of an axiom system but rather how it is used. One could take any axiom system and propose it as a foundation for mathematics or as an axiomatization of its own specific domain.

What's discussed below are a series of important properties of a particular context in which one selects between axioms, not a property of the axiomatic systems themselves.

### 1.3.1 Pure vs Impure Contexts

The first distinction is between what I call 'pure' and 'impure' contexts. Consider the following three cases involving some mathematician M and their theorem T:

Case 1: M desires to formally verify T by means of a computer-assisted proof-finder. In order to do this, they must select some axiomatic starting point that the proof-finder will work from.

Case 2: M is writing a textbook for maths students, of which an axiomatic proof of T will form a major part. They must select axioms for this textbook.

Case 3: M is in an abstract discussion with a philosopher, or especially doxastically conservative mathematician, who wishes M to justify T. M will need to select (and perhaps justify) axioms for their proof of T.

Cases 1 and 2 involve extra-mathematical considerations. As a result, they are called 'impure'. Case 1 involves considerations pertaining to the pragmatics of computer languages and available proof-finders. For instance, if a great number of the intermediate lemmas on the way to T have already been proven from some axiom system, it might be practically sensible to use those results and prove T on the same basis. In case 2, pedagogical considerations are relevant. M must consider the intelligibility of their axioms and the proofs that follow from them. It matters in that context that that proof can be easily understood.

Case 3 is what I call a pure context. Here M axiomatises simply for the purpose of doing mathematics. This is an idealised context of justification. The only considerations relevant to axiom selection in this context are those demanded by mathematics, or perhaps philosophy, in absolute abstraction.

To be clear, these purely mathematical/philosophical considerations might still be relevant in impure contexts. It's not as if, as soon as practical considerations arise, purer considerations become irrelevant. It might be the case, for instance, that there are many equivalent and objectively true axiom systems. In the pure context, there is nothing to decide between them, but in the impure contexts, the pragmatic considerations break the tie. The pure vs impure distinction should be seen as a scale, not a binary.

For the rest of this paper, I restrict the discussion to pure contexts, unless noted otherwise.

Philosophically, the pure contexts are more interesting than the impure ones. Whilst it's obviously interesting how one goes about coding, say, a proof-finder, or how one best teaches mathematics to students, there's nothing *philosophically* complex going on there. Impure contexts involve some practical end. In case 1 this is the translation of normal proofs to strict derivations. In case 2 it's teaching mathematics. Axiom selection has consequences for those ends, so can be evaluated relative to them. It's simply a matter of means-ends reasoning towards these cases.

Investigating these cases is exceptionally interesting from a computer science or pedagogical perspective, but there's nothing philosophically novel going on there.

However, the question of the right theory of axiom selection in a pure context is exceptionally philosophically interesting. It connects to a great many questions in the philosophy of mathematics about the nature of mathematical truth and knowledge, or the purpose of mathematics. As such, I focus just on the pure contexts, or the pure part of impure contexts.

One might object that, in practice, no pure context exists. Or, perhaps worse, only exists as an artefact of philosophy seminars, not mathematical practice. Practically, axiom selection always occurs with some non-mathematical considerations in mind. If this is the case, focusing on pure contexts might be uninteresting.

Even if, practically, there are no pure contexts, they are still worth attention. First, they might still be a useful theoretical context to consider, even if never practically realised. Second, as discussed above, some impure contexts might have a pure component. Understanding that pure component remains the philosophically interesting part of a study of axiom selection.

Now, It might still be that the best theory of axiom selection in pure contexts is just a null theory. Perhaps in mathematics, there are no morals and the only reasons for ever picking one axiom system over another are for practical reasons like pedagogy or proof verification. This hypothesis is very much on the table here. But that doesn't change the fact that *establishing* that this is the case is incredibly philosophically interesting and hangs on a great many open questions in the philosophy of mathematics.

In summary, there are two points to take away from this distinction. First, in practice, Axiom selection is often tied up in a number of 'impure', non-mathematical or practical matters that might have considerable weight on axiom selection in practice. Second, it's nevertheless worth isolating and focusing on the pure contexts, or the pure parts of impure contexts, as this is the more philosophically contentious, and interesting, area.

### 1.3.2 Descriptive vs Prescriptive Contexts

The second distinction, borrowing Schlimm's terminology, is between descriptive and prescriptive axiomatization. The difference between the two pertains to the relationship between the axiomatic system and the body of mathematical knowledge to which it relates. Descriptive axiomatization attempts to describe, and perhaps even explain, a body of well-defined mathematical facts. Prescriptive axiomatization fixes a body of facts. Arguably, it defines a type of mathematical object or a mathematical concept.

Something like this distinction has gone by other names in the literature. Shapiro (2005), Hellman (2003), Awodey (2004) and Fiferman's contribution to Feferman et al. (1999), for instance, have called all descriptive axiomatisation "assertory" and prescriptive "algebraic". There are some distinctions between the two, but they track essentially the same concept. Mayberry (1994) has a similar distinction between "classificatory" and "elementary" theories.

This also relates to the interpretation of the terms of an axiomatic theory, as discussed above. The prescriptive axiomatization provides what I call a 'structuralist' interpretation of the terms. The descriptive axiomatisation will provide a "Platonist" interpretation of the concepts.

As mentioned above, it's important to note that the prescriptive vs descriptive distinction is a property of how an axiomatic theory is being *used* not a property of the theory itself. One might take any axiomatic theory and, successfully or unsuccessfully, claim that it's a descriptively correct theory of some structure. Similarly, one might use the very same theory and use it to define a (structural) mathematical concept[4].

Moreover, the use of a particular axiom system might change over time. Take the history of geometry. Geometry began as a descriptive theory of the mathematical structure of real space. Euclid was engaged in a descriptive axiomatic project. However, with the discovery of non-euclidean geometry, the function of geometry changed. It ceased to be a study of real space and instead became a study of a certain kind of mathematical structure. This meant that later axiomatizations of geometry, in the knowledge of the independence of the parallel postulate, served a partly prescriptive role. They settled what kinds of mathematical structures counted as geometries. Hilbert, contra Euclid, was engaged in a prescriptive project.

As a final note, the kinds of descriptive axiomatic projects that are possible is not a philosophically neutral matter. There are two ways one might think about descriptive axiomatic projects.

[4]My comments here run contra to Shapiro and Schlimm's analyses. Schlimm, interpreting but concurring with Shapiro, states that all "assertory" (descriptive) theories need to be categorical as they are intended descriptions of particular structures. "Algebraic" (prescriptive) theories do not, as they are just defining a concept. Consequently, any assertory theory can be used in an algebraic manner but *not* vice versa. However, Shapiro is wrong that because a theory is an intended description of a fixed structure it needs to be categorical. Incomplete theories of fixed target structures abound: statistical mechanics, meteorology and the entirety of social science. Given this, there's no reason why any "algebraic" (prescriptive) theory couldn't be used in an "assertory" (descriptive) manner

First, one might think of descriptive axiomatization as an attempt to rigorize and describe a body of mathematical knowledge settled out by pre-axiomatic mathematical work. This is what Schlimm calls the ‘orthodox’ view of axiomatization.

Second, one might think of descriptive axiomatization as an attempt to accurately describe a body of objective mathematical facts. If one believes, for instance, in the objective, mind-independent existence of the natural numbers, then, one might take $PA$ to be an attempt to describe the facts about the natural numbers. The facts about $\mathbb{N}$ are already "out there" and axiomatic theories, like scientific theories, are just descriptive theories of those facts.

The possibility of the first kind of project is, I hope and suspect, relatively uncontroversial. Whatever one thinks about objectivity and mathematics, there have been bodies of pre-axiomatic mathematical knowledge that lacked, at a point in time, formal foundations. Regardless of one’s philosophical leanings, one can at least make sense of the project of systematizing and clarifying the assumptions of said project.

The second project, however, that attempts to use axiomatic systems to describe some objective mathematical facts, should be controversial. It is not an uncontroversial nor settled matter that mind-independent mathematical objects exist. Consequently, engaging in an axiomatic project that assumes their existence comes with a degree of philosophical baggage.

### 1.3.3 Local vs Global Contexts

The final distinction is between what I’ll call local and global axiomatization. This is the difference between attempting to axiomatize a specific part of mathematics (e.g. geometry, group theory, arithmetic, etc) and attempting to axiomatize mathematics as a whole (e.g. ZFC is often taken to perform this function). This latter has been referred to as a foundation for mathematics, though this is a term I try to avoid as it has quite a few varying meanings and I wish to avoid unwanted connotations.

There’s at least the theoretical possibility of an intermediate step: attempting to provide an axiomatic ‘foundation’ for some parts of mathematics, but not the whole. One might think of the axioms of, say, complex analysis as trying to provide a foundation for real analysis and discrete arithmetic.

There's some relationship between this distinction and the prescriptive vs descriptive distinction. Hellman (2003) has argued that global axiomatizations need to have "assertory content". In Schlimm's terms, they need to be descriptive, not prescriptive.

Whilst I don't think Hellman is right that 'foundations' (global axiomatizations, in my language) need to be descriptive, rather than prescriptive, Global axiomatic projects are often more natural as descriptive, rather the prescriptive, endeavours. If one attempts to provide an axiomatization for the whole of mathematics, one will typically think there is such a thing as the 'whole of mathematics' and take one's axioms to be a descriptive theory of those objects.

This isn't a hard rule. Global but prescriptive contexts exist. One might engage in a prescriptive project that defines what counts as mathematics, and hence what the entire mathematical universe is. There's a way of thinking about a certain kind of finitist project, comparable to Hilbert's program, as doing exactly this. But, typically, global axiomatization will mean descriptive axiomatization.

The converse inference doesn't hold. There are many descriptive projects that are not global. As discussed above, Euclid was engaged in a descriptive program qua geometry, but this was only local, not global.

Two further points of note on this distinction.

First, one may employ different evaluation criteria in local axiom selection compared with global axiom selection. One might have a descriptive theory of global selection (discussed below in §2) that correctly captures the entire ontology of mathematics but nevertheless have a normative theory of local axiom selection (again, see §2) that aims to find the most practical or productive distinctions within the universe of one's global theory.

Second, the adoption of a global context is not without philosophical controversy. It is a substantive assumption to make that all mathematics should adopt a shared, unifying 'foundation'.

Consider, for instance, the claim that *ZFC* should be the ultimate foundation for mathematics. To put this in realist terms, this is the claim that all non-foundational mathematics is about some sub-structure within the 'real' model of *ZFC*. To put this in more practical terms, all proofs need to take *ZFC* as their most basic assumptions. One might challenge this on two grounds. First, one might agree that there needs to be some theory that plays that ultimate unifying role but disagree that *ZFC* is the theory that should do that. Second, and this is the point I wish to stress here, one might reject the need for a unifying foundation altogether.

If one takes the latter option, then, in this paper's terminology, one is opposing a global context. One might oppose a global context in many ways. One might argue that it's impossible to take up such a context. Perhaps all possible mathematics is sufficiently broad as to make it impossible to find a single axiomatic basis for all of it. Second, one might argue against the *value* of a global context. It might be possible, but unhelpful, to argue for global axioms.

An example of someone who has arguably adopted something like this position is Awodey. He argues in favour of category theory as a 'foundation' for mathematics, though he means something quite different by the term than most others in the debate. What the category theorist is really advocating for, which I think Corfield (2020) explains quite nicely, is the use of category theory as a foundational *language* for mathematics. What different domains of mathematics then do is specify the kind of ontology they need for their particular study without any attempt to unify all of this into a single structure. There is no 'one true topos' on their view. Awodey (2004) says the following:

*"As opposed to [a] one-universe, 'global foundational' view, the categorical structural one we advocate is based instead on the idea of specifying, for a given theorem or theory, only the required or relevant degree of information or structure, the essential features of a given situation, for the purpose at hand, without assuming some ultimate knowledge, specification or determination of the 'objects' involved." (p56.)*

What's clear, then, is that the necessity of axiom selection for a global context is not philosophically uncontroversial.

### 1.4 Summary

In summary of the preliminaries, axiomatic systems comprise of a set of statements in some language, a consequence relation and (possibly) an interpretation of the non-logical symbols. Axiom selection always takes place within a particular context. Theories of axiom selection are correct modulo some context. The context might be pure and only have mathematical (or perhaps philosophical) ends, or it might be mixed and contain certain impure, extra-mathematical ends. This paper focuses on the former. Axiom selection might relate to a body of mathematical truths in one of two ways. First, it might define said truths and be prescriptive. Second, it might try to describe a pre-determined body of truths. Lastly, axiom selection can happen at various levels of generality. Sometimes axioms are selected for singular parts of mathematics (e.g. the axiomatisation of a particular domain) and sometimes axioms are selected as part of a global or foundational project, though the necessity or value of such is not a philosophically neutral matter.

## 2 Evaluation criteria for axiom systems

Having fixed the context in which axiom selection takes place, a theory of axiom selection must explain the evaluation criteria to be employed to choose between rival systems. The kinds of criteria available are quite diverse and there is, I suspect, no philosophically uncontroversial 'common ground' on which to select one's evaluation criteria.

I do not argue here for any particular criteria over others. This section's aim is simply to provide a taxonomy of the kinds of criteria available and the kinds of relations that might hold between them. I group criteria into four types: descriptive, epistemic, conceptual and normative. Each are explained in turn.

An important note here is that even after fixing the evaluation criteria, one still does not have a complete theory of axiom selection. There's a gap between the claims that "T is the best axiom system (in context C)" and "One ought select T and only T for use (in C)". The gap exists on two grounds. First, one needs to establish that one can only select a single axiom system. At least theoretically, there's the possibility of plural foundations. Second, one needs to establish that some axiom system needs to be selected *at all*. In other words, the 'null foundation' needs to be ruled out.

The second part of the gap can be handled relatively easily by including the 'null foundation' as a possible axiom system to be evaluated. It's then clear if some candidate axiom system is, in fact, better than nothing.

The possibility of a plural foundation, however, requires further discussion. Axiomatic Pluralism, as it could be called, is a position I defend elsewhere. §3 discusses three types of Axiomatic Pluralism, again aiming to categorise them, rather than defend them. The question of pluralism is put aside till then.

## 2.1 Descriptive Evaluation Criteria

Descriptive evaluation criteria treat the relationship between axiom systems and mathematical facts as analogous to the relationship between scientific theories and scientific facts. The term analogous should not be here-understood as too strict a relationship; there may be differences between the two. But, in essence, the two relationships are treated similarly.

Examples of descriptive criteria on axiom selection include ontic conditions discussed by Ladyman and Presnell (2016), and the classical view's requirement that axioms be true, also reiterated in Mayberry (1994). Arguably, Hellman (2003)'s requirement that axiom systems provide 'assertory content' also falls in this category, though there are potentially some interpretive questions there.

A standard realist claim about scientific theories is that they should be true generalisations of the scientific facts. For instance, Hooke's Law, if it is to be a correct theory of springs, should entail all and only the facts about springs.

This idea can be translated quite easily into a claim about axiom systems and mathematical truth. An axiom system, if it is to be correct (or correct of some domain), should entail all and only the mathematical facts (of that domain).

Something like this is, I think, a very natural idea to have of axioms. They are, at least logically, the foundational truths of mathematics from which all further knowledge is derived.

Of course, the requirement that an axiom system should entail all the truths of mathematics runs into immediate difficulties with Gödel's first incompleteness theorem. The first incompleteness theorem entails that, modulo some facts about encoding arithmetic, there are no finite, consistent, complete first-order, classical theories.

This potentially necessitates a slight weakening of the requirement that axiom systems entail all and only the mathematical truths.

First, as discussed in §1.1, it might be necessary to allow for infinitary axiom systems (in the sense of systems comprised of more than finitely many axioms and axiom schema). It might then be that there are human limitations on our ability to ever completely know what these systems are, but they would, nevertheless, be the correct theory.

Even if this is the case and the completely correct axiom system is infantry, it's still useful to introduce the notion of an axiom system being more or less correct than another. Even if there is no finitary axiom system that can entail all the mathematical truths, some might nevertheless entail more, or more important, truths than others.

A second option is to adopt an axiom system with a non-classical logic. This could either be systems with higher-order axioms or perhaps with a non-classical entailment relation. If one took the former route, there's a question as to which notion of entailment is relevant for the evaluation of axiom systems: semantic ($\models$) or syntactic (). Higher-order theories can be categorical but are not logically complete.

Returning to the question of descriptive evaluation criteria, there are other requirements one might place on laws of nature beyond descriptive accuracy. Often, theories of laws have a nomological requirement. Laws do not merely describe the natural facts, they *explain* them (Maudlin, 2007).

A similar condition could be put on axiom systems. Axioms should not merely describe the mathematical facts but explain them. Unfortunately, it is an open question what the correct notion of explanation within mathematics is, and even if there is one at all. This is not to be confused with mathematical explanation in science, about which there is perhaps more understanding but less agreement. See D'Alessandro (2019) for an overview and Lange (2014) for an influential recent work.

Elsewhere, Lange (2018) discusses the idea that parts of mathematics other than proofs might explain. A natural option, though not one discussed by Lange, is that axioms might explain certain mathematical facts. D?Alessandro (2018) considers something close to this, though ultimately rejects "inter-level" explanations in mathematics. If the idea of explanation in mathematics can be refined and extended to the relationship between axioms and theorems, an explanatory criterion on axiom selection could prove attractive.

One advantage of including an explanatory condition on axiom selection is that it would allow one to choose between rival logically equivalent systems.

Choosing between rival systems might not be something strictly necessary. It might not be the case, for instance, that one needs to work out which of the Axiom of Choice and the Well Ordering Principle is really correct! Having to make such judgements seems a little silly.

However, one could imagine finding a number of intermediate lemmas of *ZFC* that taken together, it just so happens, are equivalent to *ZFC*. These lemmas are all a little arbitrary and clearly less 'primitive', in a certain sense, than the axioms of *ZFC*. It would likely be useful to have a way of ruling such theories out. Mere descriptive criteria cannot do this. Explanatory criteria, however, could, as explanation is more fine-grained than logical equivalence.

As a final note in this section, descriptive criteria will normally be adopted only by mathematical realists. If axioms need to be true, or explanatory, or something else which entails truth, there need to be mathematical truths. That, or one needs to accept axiomatic nihilism, rejecting all axiomatic theories. But I take that to be an extreme that no one is actually willing to take up. Baring this extreme, though, only realists can adopt

Now, I don't necessarily see this as a negative feature of descriptive criteria. If one is a realist, then one's realism should matter for mathematics, likely via mattering for axiom selection. But it should be understood that their realism is an assumption being made by employing descriptive evaluation criteria in axiom selection. Descriptive criteria should not be uncritically presented as theory-neutral.

## 2.2 Epistemic Evaluation Criteria

One of the most distinctive features of the classical concept of an axiom (or axiom system) is self evidency. A reasonably well-known picture of mathematical knowledge is typically outlined along with this. In mathematics, one starts with knowledge of the axioms. These are statements that, upon reflection or understanding of the concepts involved, are known for certain to be true. From these sure foundations, the rest of mathematical knowledge is derived.

The person in recent years to have advanced a picture at all like this is Mayberry (1994), though even he admits that complete self-evident certainty is not a reasonable requirement on axiom systems. Short of self-evidency, a common idea is that axioms should play some kind of justificatory role in relation to the rest of mathematical knowledge. Linnebo and Pettigrew (2011) talk of justificatory autonomy: foundational axiom systems should be justified without reference to other mathematical knowledge. Ladyman and Presnell (2016) discuss a number of possible epistemic requirements, though don't actively endorse any specific claim in this regard.

Epistemic conditions can betake two forms. They can be about the level of justification required of axioms, but they can also be about the epistemic link between axioms and mathematical knowledge (call this 'link strength').

The level of certainty required of axioms is a relatively easy-to-understand variable. Some epistemic conditions might require that axiom systems be known with certainty, others that they merely be known. Within the latter, there's a question about what the bar of justification is. Different epistemic conditions might place higher and lower bars.

Tied into this is not merely the strength but the *kind* of justification of the axioms allowed. The classical view requires not just that axioms are known but they are known self-evidently. Linnebo & Pettigrew's justificatory autonomy condition requires that axioms are justified (perhaps not with certainty) without dependence on other areas of mathematics.

Epistemic conditions can therefore demand that axiom systems not only be justified to a certain degree, but justified *in a certain way* and to a certain degree.

Moving on to link-strength, There are many possible views about the relationship between axiomatic knowledge and normal mathematical knowledge to be. One might hold, in decreasing strength, that: (1) axiomatic knowledge is essential for ordinary mathematical knowledge. Only via derivations from the axioms can a mathematical claim be known. Arguably, Mayberry (1994) has this view. (2) Axiomatic knowledge contributes towards ordinary mathematical knowledge, but not strictly necessary. It's possible to have mathematical knowledge without axiomatic knowledge, but perhaps more difficult. Lastly, (3) axiomatisation is irrelevant, or only minimally relevant, to mathematical knowledge.

(1) and (2) lead to natural epistemic conditions on axiom selection. A proponent of (1) might hold that one ought evaluate axiom systems on their ability to provide this essential foundation. A proponent of (2) might hold the weaker claim that axiom systems should *contribute* towards mathematical knowledge, but it's possible to have mathematical knowledge without it.

(3) might be held by individuals like Copi (1958), Weyl (1935) or Hanson (1958a), as discussed in Schlimm (2013), who hold that axiomatisation merely involves the systematisation of known mathematics, not the generation of any new knowledge.

As a counterpoint, one might also consider the reverse direction in which axioms are epistemically dependent upon ordinary mathematical knowledge. Schlimm, Maddy (2011) and Quine (1969) all defend views whereby axioms can be justified on the basis that they either entail or explain known mathematics. In Schlimm especially, a cyclical feedback loop is emphasised. Some ordinary mathematical knowledge justifies axiomatic knowledge, which in turn justifies more ordinary mathematical knowledge.

This cyclical conception of the relationship between ordinary and axiomatic mathematical knowledge is antithetical to stronger epistemic conditions on axiom selection. The cyclical view is opposed to the view of axioms as an epistemic foundation for mathematics.

In summary, one might evaluate axiom systems on the basis of the following kinds of epistemic conditions. One might require that they themselves are justified with a certain degree of warrant, perhaps even restricting one's self to certain kinds of justification. Moreover, one might require that axioms are, in a certain sense, epistemically prior to ordinary mathematical knowledge. Axiom systems might then be essential to or contribute towards ordinary mathematical knowledge.

Two points of note here.

First, as with descriptive criteria, epistemic criteria are built on realist assumptions about mathematics. Knowledge is typically understood as factive, it entails truth. So in order for there to be knowledge there needs to be mathematical truths. This is a claim that only a realist is entitled to[5].

As with descriptive criteria, I don't per se take that to be a positive or a negative feature of epistemic criteria on axiom selection. It is, again as with descriptive criteria, an assumption that proponents of epistemic criteria should be open about and they should not be presented as theory-neutral.

Second, to my knowledge, epistemic criteria have only ever been applied to axiom selection in a global (i.e. foundational) context. Some epistemic criteria would make little sense in a local context, e.g. Linnebo & Pettigrew's epistemic autonomy condition. Other criteria might make sense. One might require that the axioms of, say, geometry play a certain role in the justification of geometric knowledge.

## 2.3 Conceptual Evaluation Criteria

The penultimate type of criteria for axiom selection is conceptual criteria. Conceptual criteria evaluate axiom systems on the basis of their ability to define or aid in the understanding of mathematical concepts.

[5]Of course, a non-realist is welcome to adopt both a non-factive account of knowledge and epistemic criteria on axiom selection. But such accounts of knowledge are few and far between.

There's a structural parallel between the kinds of epistemic and conceptual conditions that one might adopt. One might evaluate the concepts present in the axioms themselves. For instance, Maddy's "intrinsic" considerations or Linnebo & Pettigrew's conceptual-autonomy condition. And, again as with epistemic conditions, one might consider the relationship between axiom systems and axiomatic concepts (i.e. the concepts present in the axioms) and 'ordinary' mathematical concepts. Both Mayberry, and Ladyman and Presnell have the relatively minimal requirement that axiom systems should be able to define ordinary mathematical concepts. For instance, many logical notions can be easily defined in homotopy type theory.

Of axiomatic concepts themselves, one might require a number of things. Maddy takes it to be a virtue of certain axioms if they are the natural explication of certain, presumably pre-axiomatic, mathematical concepts; so-called 'intrinsic consideration'[6]. Linnebo  Pettigrew's conceptual autonomy condition evaluates axiom systems based on if their concepts can be defined or understood independently of other mathematical concepts. Though, as with the epistemic case above, this doesn't seem to be a criterion applicable to local contexts.

A crucial philosophical question here, returning to the discussion §1.1, is if a purely logical or 'structural' understanding of axiomatic concepts is sufficient for these kinds of criteria. Is it sufficient in order to understand some mathematical concept that one merely knows the kinds of logical inferences one can make from it (see discussions on inferentialist semantics in mathematics)? Or does one need some richer, more philosophically loaded kind of understanding? Linnebo  Pettigrew do consider a stronger notion of understanding than mere 'structural' understanding, though they don't necessarily endorse such a requirement. One might endorse conceptual criteria on axiom selection that require either.

[6]It is worth noting that she weighs intrinsic considerations far less heavily than extrinsic considerations (see discussion in §2.4). There are passages later in the book where she can even be read as dismissive of the value of intrinsic considerations altogether.

Moving on to the relationship between axiomatic concepts and ordinary concepts, the same kinds of things can be said as with the epistemic case. One might require quite a strong link between one's axiomatic concepts and one's ordinary concepts whereby the understanding of axiomatic concepts is entirely necessary for understanding ordinary concepts. This would be, I think, really quite a strong condition and it has not to my knowledge been defended. More likely, one would require that one can get understanding of ordinary concepts via one's axiomatic concepts.

Most conceptual criteria around the link between axiomatic and ordinary concepts that have been defended in the literature take an even weaker position still (Mayberry (2000), Ladyman and Presnell (2016)). One only needs formally equivalent *surrogates* of ordinary mathematical concepts to be definable from one's axiomatic concepts. $\{x\} \cup x$ need not literally be the successor relation on the ordinals. It's sufficient that it captures the requisite formal structure.

I'll note the possibility that this entire section is a red herring. I think it's entirely possible that conceptual criteria shouldn't really be seen as criteria in their own right, but rather as necessary conditions for other criteria. In Mayberry (2000), for instance, it's clear that definitional power is only valuable because one needs to be able to do mathematics from one's axioms. If you can't define (surrogates of) ordinary mathematical concepts, you can't do ordinary mathematics with your axioms. If one endorses basically any descriptive, epistemic or instrumental (see §2.4 below) criteria, one's probably committed to some relatively weak conceptual criterion as a result.

## 2.4 Normative Evaluation Criteria

As noted above, descriptive and epistemic evaluation criteria come with realist assumptions. What I call 'normative' evaluation criteria do not. They are intended as value-based evaluations of axioms that don't attempt to justify axiom systems on the basis of their relation to any mathematical facts. Descriptive and epistemic criteria are concerned with finding the *correct* axioms, normative criteria are concerned with finding the *best*.

That is somewhat of an open description because this is somewhat of an open category. There are, I think, two sensible entries under this category: instrumental considerations and simplicity. These are discussed below. However, there are a great many more considerations that, whilst implausible, are certainly logical possibilities. Berkeley, for instance, opposed algebra on theological grounds (Moriarty, 2018). Theological considerations would certainly count as normative as normative criteria. That a certain axiom is, say, beloved by God needn't mean that God made it true. Another example might be aesthetic considerations. Perhaps the beauty of a particular axiomatic theory might count in its favour. Both of these considerations likely aren't *good* normative criteria, though this is where they would fit within my conceptual framework if one were to defend them.

The best examples of an instrumental consideration in favour of some axiom system are Maddy (2011)'s "extrinsic" reasons for axiom selection[7]. For Maddy, the main basis for selecting one axiom over its rivals is its utility for mathematics. Axioms need to entail the existence of interesting or deep mathematical structures. They need to populate a mathematical universe with the kinds of objects worth mathematical consideration. One can adjudicate between rival axioms based on their ability to perform this task.

I discuss Maddy's view extensively in Essay III: *Instrumentalism & Axiom Selection* but a few quick comments are nevertheless in order here.

First, on Maddy's view, rival axiom systems are to be judged on the basis of their utility for mathematics. A natural variant of this view, one that I defend in the aforementioned paper on Maddy, is to instead judge axioms on their utility for science. Science needs modelling tools. Mathematics provides some of the best formal modelling tools. Axiom systems should be judged on their ability to provide said tools. More generally, one might consider the utility of mathematics for all kinds of activities that make use of mathematics, science and mathematics itself are just the most natural examples.

[7] See also Easwaran (2008)'s account.

Second, an important consideration left out of Maddy's account is the possibility of axiom pluralism. In order to show that one ought adopt some axiom A, it's insufficient to show that A has more utility than its negation. If both A and its negation entail the existence of interesting or fruitful structures, a pluralist approach that investigates both the consequences of A and the consequences of not A is likely more appropriate. I discuss this at length elsewhere.

Lastly, Maddy really only talks about the adoption of axioms in a global context. She doesn't consider local contexts. Nevertheless, her view generalises quite easily to those contexts. The only difference is that whereas axioms in a global context are judged on their ability to deliver a wide range of interesting mathematical structures; a sandbox for all of mathematics, axioms in local contexts are judged on their ability to produce a single interesting, deep or fruitful type of structure.

To return to the metaphor of mathematics as a tool, in a global context the instrumentalist wants an axiom system to be a good toolbox. It will be judged on both the breadth and quality of its tools. In a local context, the instrumentalist wants an axiom system to be like a good singular tool. Only the quality of that one tool matters.

The second type of normative criterion I think is worth specifically mentioning is a certain kind of simplicity-to-depth ratio that is of some kind of mathematical value. Interesting mathematics is often characterised not merely by its depth but by its depth despite the apparent simplicity of the definitions or axioms on which it rests. Maddy has discussed this at some length. Group theory is, in my opinion, one of the nicest examples of this. The definition of a group can be understood by high school students; it's a very simple algebraic concept. However, such a simple notion has a wide range of interesting and powerful applications, from describing certain kinds of geometric transformations to conservation laws in mechanics.

Now, properly understood, simplicity should never come at the cost of depth. It is likely better to have bloated but effective axioms than slim but dysfunctional ones. However, all else being equal, a more elegant, simple axiom system is a better one. This analysis follows Schlimm (2013)'s.

As a final note on normative conditions, it's worth considering for a moment the kinds of 'impure' contexts ruled out in §1.2.1. Largely, I ignore these contexts in this paper. However, it's worth briefly mentioning that the kinds of considerations that become relevant in impure contexts would typically be classified as normative criteria. Pedagogical value, for instance, would fall here.

## 2.5 A conscious omission: Logical Criteria

Schlimm discusses a range of criteria for axiom selection which he groups under the heading "Meta-mathematical Properties" (§4.3). Criteria such as mutual independence, consistency and categoricity have all been used as desiderata when deciding between axiomatic theories. Those familiar with the history of axiomatics might have been surprised by the absence of a discussion of categoricity or consistency above, given their ubiquitous historical importance (see Schlimm for discussion). One might expect them to be obvious candidate criteria for axiom selection.

However, I think these kinds of logical criteria are typically (perhaps even universally) advanced as secondary criteria, derived from other criteria of axiom selection, not as basic considerations themselves. It's clear to see the kinds of considerations that might lead one to value the aforementioned logical considerations. If one adopts descriptive criteria, the need for consistent theories follows. Inconsistent theories have no models so cannot describe any structures. Similarly, there are descriptive, conceptual and normative reasons to wish to unique categorise structures. This would allow for a complete description of some real mathematical structure (descriptive), a maximally precise concept (conceptual) or a fine fine-grained and precise representational tool (normative). Categoricity would then be a derivatively valuable property of axiom systems.

Dedekind (1888), for instance, took the concept of an arithmetical structure to be precisely the structure $\mathbb{N}$. An axiomatisation of arithmetic, in Dedekind's eyes, would only be successful if it could uniquely characterise that structure. In other words, if all of its models were isomorphic to $\mathbb{N}$. If it was categorical.

Crucially, then, categoricity is not a fundamental criterion of axiom selection for Dedekind. Instead, he employs a conceptual criterion: his desire for a unique characterisation of a particular mathematical concept, in this case, the concept of an arithmetical structure. From this conceptual criterion, she derives the categoricity condition.

I believe that Dedekind's case is typical for historical examples of logical criteria on axiom selection. Logical criteria are infrequently defended as fundamentally valuable but frequently defended as *consequences* of other criteria for axiom selection. See Awodey and Reck (2002a) and Awodey and Reck (2002b) for a more thorough discussion of the historical cases. For the sake of brevity, I will simply state my claim that, as in Dedikin's case, logical criteria on axiom selection like categoricity are infrequently if ever *fundamental* criteria on axiom selection. A more thorough analysis of the history of categoricity which demonstrates this claim is left for future work, though readers are invited to check Awodey & Reck's papers to check this claim for themselves.

Consequently, I do not include logical criteria as types of evaluation criteria for axiom selection, despite their obvious historical importance. I see logical criteria as consequences of other evaluation criteria, not criteria in their own right. They are therefore omitted.

This does, however, raise the question of exactly how these various logical criteria are to be motivated, and on what basis. There has been a great deal of recent work on categoricity and its philosophical consequences. For instance Meadows (2013), Button and Walsh (2018) and Maddy and Väänänen (2023). I hope that an analysis on the basis of this paper's framework for thinking about theories of axiom selection would be productive. It is a question I hope to return to in future work, though for the sake of scope must be left unanswered in the rest of this paper.

## 2.6 Summary

In summary, the following is my 'conceptual map' of the possible evaluation criteria used in axiom selection:

1. Descriptive Evaluation Criteria
    - Merely descriptive criteria:
        - How should the issue of incompleteness be handled?
        - In logics where they come apart, which of categoricity or completeness should be favoured?
    - Explanatory criteria
        - What's the relevant notion of mathematical explanation here?
2. Epistemic Criteria
    - Qua the axioms themselves:
        - How well justified do the axioms need to be, if at all?
        - What kinds of justification are permitted in justifying the axioms?
        - What's the relationship between axiomatic knowledge and 'ordinary' mathematical knowledge?
    - Qua the link between axiomatic knowledge and ordinary mathematical knowledge:
        - Do axioms need to contribute towards ordinary mathematical knowledge?
        - If so, is this necessary or merely contributory?
3. Conceptual Evaluation Criteria:
    - Qua the axioms themselves:
        - Do we need to understand the concepts used in our axioms?
        - Is a mere 'structural' understanding good enough?
        - Can we rely on other mathematics in getting this understanding?
    - Qua the link between axiomatic concepts and ordinary mathematical concepts:
        - Do axiomatic concepts need to contribute towards our understanding of ordinary mathematical concepts?

    - If so, is this necessary or merely contributory?
    - Are formally equivalent surrogates of ordinary concepts sufficient?

4. Normative Evaluation Criteria:

    - Instrumental Criteria
        - Instrumental value for maths, science or something else?
        - Quality vs quantity of mathematical structures
    - Simplicity and depth
    - Other normative criteria not discussed (pedagogical value, aesthetic considerations, etc)

These criteria are not exclusive. One is able to combine them in interesting and novel ways, varying their relative weight and priority. A theory of axiom selection should, for any particular context, identify which of these criteria are relevant for evaluating one's axioms, and how to weigh competing criteria against one another.

My hope is that this section clarifies two things. First, the kinds and the diversity of possible evaluation criteria that one could use. The present literature, I hope it's clear, has not mapped all the possibilities. Second, there are substantial and controversial philosophical assumptions behind one's choice of evaluation criteria and, consequently, one needs to be open and explicit about one's theory of axiom selection *and the philosophical assumptions behind it*.

# 3 Axiomatic Pluralism

The last section of this paper discusses a natural corollary of several points mentioned in both §1 and §2: the possibility of axiom pluralism. At several points in this paper, I indicate opportunities for possible pluralist positions about axiom systems. In this section, all of that is brought together and a number of possible versions of "axiomatic pluralism" are outlined.

Axiomatic pluralism (of this sort) is a position that has been identified but not seriously within the present literature. Michèle Friend has conducted the most extensive research on pluralism in mathematics generally. She discusses "foundational pluralism", but not in extensive depth (Friend (2013) and Friend (2019)). Davies (2005), Hellman and Bell (2006), Koellner (2009) and Priest (2019) have all discussed the application of *logical* pluralism to mathematics, but this is a very particular kind of axiomatic pluralism and not representative of the position at large.

What exactly is "axiomatic pluralism"?

In response to the question "How many correct logics are there?" one must give one of three answers: none (nihilism), one (monism) or many (pluralism). Logical pluralism is the view that there are many correct logics. More generally, when asked "How many correct theories of X are there?" one can, again, answer none, one or many and be an X-nihilist, X-monist or X-pluralist respectively.

It's natural to understand axiomatic pluralism in the same way. In response to the question "How many correct axiom systems are there?" one might answer none, one or many and be a nihilist, monist or pluralist respectively. Axiomatic pluralism is therefore the view that there are many correct axiom systems.

Two comments here.

First, the term "correct" potentially has some realist connotations that I wish to avoid. As is shown above, one needn't have a realist theory of axiom selection that requires axioms to be *descriptively* correct. I certainly don't intend to bake realist assumptions into the notion of axiomatic pluralism.

The notion of "correct" at play here should be read neutrally. It might be read as descriptively correct, if one's background theory of axiom selection permits, but it might also be read as *normatively* correct, in the sense of a correct choice or course of action, if more appropriate.

Second, there's a comparatively weak pluralism that it might be helpful to rule out. This is briefly discussed in §2.1. It might be that one's theory of axiom selection is unable to decide between logically equivalent axiom systems. There might be, for instance, no grounds to prefer Choice to Zorn's lemma as an axiom, or vice versa. This kind of axiomatic pluralism would be comparable to a kind of weak logical pluralism that is pluralist on the grounds that one can't decide between, say, different but equivalent sets of inference rules or different basic connectives.

Now there have been logical monists who have argued that there are facts of the matter about which connectives are really the basic ones (McSweeney, 2019). So the option of defending the analogous position in the philosophy of mathematics should be seen as a live one. I certainly don't want to take weak pluralism as a given. However, if one's 'pluralism' only goes as deep as endorsing multiple logically equivalent axiom systems, one isn't a pluralist in any interesting sense. If such a position should even count as pluralist at all I leave to the reader. It would simply be a matter of amending the notion of "axiomatic pluralism" to mean pluralism about the number of equivalence classes of correct axiom systems, rather than pluralism about the systems themselves.

There are three types of axiomatic pluralism outlined in this section: multiverse pluralism (I take the name from the set-theoretic multiverse project), domain-relative pluralism (à la Pedersen (2014)'s domain-relative logical pluralism. See Shapiro and Lynch (2019) for a nearby discussion. Criticised by Steinberger (2019)) and theory pluralism.

The aim here is not to defend any of these kinds of pluralism, merely to identify them for future work.

## 3.1 Multiverse Pluralism

The first type of pluralism I call 'multiverse pluralism'. The name is drawn from a well-known project in set theory, and philosophy of set theory, the multiverse program (Steel's contribution to Feferman et al. (1999), Hamkins (2012), Steel (2014), Antos et al. (2015), Maddy and Meadows (2020)). The core philosophical idea behind the multiverse program is the claim that set theory is no longer, and should not be, about a singular set-theoretic universe $\mathcal{V}$. Instead, Set Theory is about a family of possible models of various set-theoretic axioms; although, models of *ZF* and its extensions form the backbone, if not the entirety of this family. Unlike, arguably, arithmetic, Set Theory has no intended model.

Now one should be careful here. Adopting a multiverse view of set theory doesn't necessarily entail axiomatic pluralism. There have been many attempts to provide multiverse axioms. In such case, it might be better to describe the resulting position as axiomatic *monism* with multiverse axioms, rather than pluralism. Personally, I'd still see that as a pluralist view, given that the axioms just give a systematic way of describing a plurality of set-theoretic universes. However, there's some reasonable disagreement to be had there.

It is, however, irrelevant. The point is that there are possible versions of the multiverse program that achieve their ends simply by adopting a kind of axiomatic pluralism. Instead of taking *ZF* or *ZFC* to be the axioms of set theory, a multiverse pluralist instead holds that there are many different axiom systems for set theory. Set theorists are interested in the consequences of many axiom systems, not merely one.

Nothing said in the above paragraphs turns on the details of set theory as the case in hand. One might just as easily say essentially the same thing about any area of mathematics. In fact, essentially the same thing has been said in relation to Homotopy Type Theory (Corfield (2020), Awodey (2004)), though other presentations differ (Linnebo and Pettigrew (2011) focus on the Elementary Theory of the Category of Sets as a universe-like foundation for mathematics). More generally, about any domain of mathematics one might reject the existence of a privileged or intended model in favour of a multiverse view, whereby the domain considers a class of possible models and a class of possible axiom systems characterizing these models. For this domain, one could then be an axiomatic pluralist.

The most famous example of a multiverse in mathematics, though it has not gone by this name, is geometry. Since the discovery of non-Euclidean geometries, geometers have not been interested in a single intended model of geometry. Instead, they are interested in a range of models and the kinds of properties those models might have.

Multiverse Pluralism generalises easily to mathematics as a whole. One simply holds that for mathematics in general, there is no privileged or intended model and therefore no singular correct axiom system. Mathematicians are not or should not interested in the consequences of one particular axiom system, but in the consequences of a range of systems instead.

Now, a natural reply here by the axiomatic monist would be to take a supervaluational approach: to claim, perhaps, that the correct axioms for a multiverse-domain are some systematization of the statements true of all the models and that statements true of only some models are to be understood as parts of conditionals.

This is certainly an interpretation of mathematical domains without intended models. Its relative merits and virtues compared to the pluralist interpretation shan't be expanded upon here. As stated, my aim here is simply to sketch the kinds of pluralist positions available, not to argue for them.

I will note, however, that the disagreement in part turns on an issue of the correct theory of axiom selection. The monst, here, takes supervaluational truth to be the relevant notion of truth for in axiom selection. The pluralist believes that truth-in-a-model is more important.

## 3.2 Domain Relative Pluralism

The second type of axiomatic pluralism is comparable to domain-relative logical pluralism, as defended by Pedersen (2014) and criticized by Steinberger (2019).

Domain relative logical pluralism holds that the correct logic can vary from domain of discourse to domain of discourse. There might be different correct logics for fundamental physics, social science, art, ethics, etc. The view is pluralist because there are multiple logics correct for some domain.

Interestingly, this is also a claim that some logical monists agree with. Priest (Priest (2001), Priest (2005)), for instance, thinks that the correct logic might vary from domain to domain. However, Priest is a logical monist because he takes there to be a 'one true logic' that's the intersection of the correct logics for each domain. A natural way to think of this is via a mereology of domains, with the general domain being the universe, in the mereological sense. The one true logic is the logic of the universe.

What this means is that for one to be a domain-relative logical pluralist, at least in any interesting sense, one has to hold not just that the correct logic can vary from domain to domain, but that there's no privileged or universal context whose correct logic would be privileged.

How does this relate to axiomatic pluralism?

Now, there's clearly no disagreement that different mathematical domains, or different local contexts in this paper's terms, have different correct axiom systems. The correct axiomatization of geometry is different from the correct axiomatization of arithmetic, is different from the correct axiomatization of set theory, etc. This is something that the monist agrees to, they simply think that there's a privileged domain (the global context) whose correct axioms are the really correct axioms.

Returning to the mereological idea, it's quite natural to think of mereological relations between mathematical domains. Discrete analysis is, perhaps quite literally, part of real analysis, which is part of complex analysis. Perhaps all of that is part of set theory, still. The monist then believes that there's a universal domain, a general context in this paper's terms. The properly correct axiom system is the system for the general domain.

To clarify, I don't want to overstate the monist's commitments here. There's a very literal reading of the last section whereby numbers, say, are *quite literally* sets. To my knowledge, only one person has actually defended that quite extreme position (Steinhart, 2002). A more common position is that structurally equivalent surrogates of, say, the natural numbers can be found in models of *ZFC*. The existence of a general domain is then just the existence of a domain in which surrogates of all mathematical structures can be found, which acts as a kind of court of appeals for issues not resolvable in the lower-level theory.

As discussed in §1.2.3, there's a debate to be had as to if there is or should be a single unifying domain for mathematics. To put this in more pragmatic terms, it's not a philosophically neutral matter as to if mathematics *needs* a single unifying, general context.

The domain-relative pluralist simply rejects the need for a universal, general domain. Of course, they acknowledge more or less 'general' domains, in the sense that there are formally stronger and weaker theories that contain surrogates of other theories. But they see stronger theories not necessarily as rivals, but as multiple possible tools all of which can be used to do more abstract mathematics. They reject the need for *a* foundation, in favour of a plurality and hierarchy of theories able to do foundational work.

The domain-relative pluralist, for instance, sees no reason to choose between homotopy type theory (or, rather, specific categories described in the *language* of HoTT) and, say, *ZFC*. They see both of these as viable, useful and co-tenable abstractions of their lower-order structures. Crucially, though, they take neither of these to be the or even a universal domain. Both are simply *other* domains, which happen to contain surrogates of logically weaker theories, not prioritised domains.

## 3.3 Theory Pluralism

The final type of axiomatic pluralism is theory pluralism. The theory pluralist believes there are multiple correct theories of axiom selection, perhaps each correct relative to some view of the purpose or function of mathematics. Perhaps, for instance, one might hold that there is no such thing as axiom selection in abstract, but rather always for a particular task. What's required of axioms might differ greatly when picking axioms for automated proof checking, than when using them to teach mathematics.

The monist might accept all this but reply that there's nevertheless a privileged function of axioms. Perhaps the task of truly *justifying* mathematics, or some such similar highly-minded task. The theory pluralist simply rejects that there is any such context that can be seen as privileged in the relevant way. I suspect this will typically couple with a more practical and less philosophical view of what axioms do. Certainly, if axioms are needed to justify mathematics, that seems sufficient for

A potentially helpful comparison here, to return to examples from logical pluralism, is Beall and Restall (2005)'s consequence-pluralism. Beal & Restall grant that, given a fixed consequence relation, there is a singularly correct logic. Similarly, they don't think that logical correctness varies from context to context. Nevertheless, there is no privileged consequence relation, there are many. Consequently, there are many correct logics; each correct modulo some consequence relation.

Just as Beall & Restall grant that for each consequence relation there is a single correct logic, but that there is no privileged consequence relation, so too does the Theory Pluralist grant that there might be a single correct axiomatic theory for each set of evaluation criteria, but there is no single privileged set of evaluation criteria.

A natural way to arrive at theory pluralism is via pluralism about the function of axiomatisation. If axiomatisation has many functions, some pedagogical, some proof-theoretic, some philosophical, then there might be multiple correct theories of axiom selection, modulo each of these functions. These different theories might endorse different axiom systems. Consequently, there would be many many correct axiom systems and one would be an axiomatic pluralist.

Returning to my previous distinction between pure and impure contexts, the Theory Pluralist might, with good cause, doubt the value of this distinction. Trying to pick axioms devoid of their practical use in automated theorem checking or proving, teaching, etc might reasonably be criticised for detaching axioms from a number of the contexts in which they are so important. The Theory Pluralist would very likely resist my previous (practical) decision to restrict the discussion only to pure contexts.

# 4 Conclusion

This paper aims to fill a present gap, or perhaps partial gap, in the literature on axiom selection. It aims to provide a *conceptual geography* of the kinds of theories of axiom selection one might take up. My hope is that this map outlines and clarifies some of the possible points of disagreement about axioms and how to evaluate them. Following from this, I hope that it's clear that axiom selection is not a philosophically neutral matter. One's theory of axiom selection will have to make contentious claims in the philosophy of mathematics; claims about mathematical truth, ontology, knowledge, and the value or function of mathematics. There is, I believe, a rich philosophy of axiom selection not yet properly explored or understood.

## References

Antos, C., Friedman, S., Honzik, R., and Ternullo, C. (2015). Multiverse conceptions in set theory. *Synthese*, 192(8):2463–2488.

Awodey, S. (2004). An answer to hellman's question: Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 12(1):54–64.

Awodey, S. and Reck, E. H. (2002a). Completeness and categoricity. part i: Nineteenth-century axiomatics to twentieth-century metalogic. *History and Philosophy of Logic*, 23(1):1–30.

Awodey, S. and Reck, E. H. (2002b). Completeness and categoricity, part ii: Twentieth-century metalogic to twenty-first-century semantics. *History and Philosophy of Logic*, 23(2):77–94.

Beall, J. and Restall, G. (2005). *Logical Pluralism*. Oxford, England: Oxford University Press.

Button, T. and Walsh, S. P. (2018). *Philosophy and Model Theory*. Oxford, UK: Oxford University Press.

Copi, I. M. (1958). Artificial languages. In Henle, P., editor, *Language, Thought, and Culture*, chapter 4, page 96–120. Ann Arbor: The University of Michigan Press.

Corcoran, J. (1973). Gaps between logical theory and mathematical practice. In Bunge, M. A., editor, *The Methodological Unity of Science*, pages 23–50. Boston: Reidel.

Corfield, D. (2020). *Modal Homotopy Type Theory: The Prospect of a New Logic for Philosophy*. Oxford, England: Oxford University Press.

D?Alessandro, W. (2018). Arithmetic, set theory, reduction and explanation. *Synthese*, 195(11):5059–5089.

D'Alessandro, W. (2019). Explanation in mathematics: Proofs and practice. *Philosophy Compass*, 14(11):e12629.

Davies, E. B. (2005). A defence of mathematical pluralism. *Philosophia Mathematica*, 13(3):252–276.

Dedekind, R. (1888). *Was sind und was sollen die Zahlen?* Verlag von Friedrich Vieweg und Sohn, Braunschweig, 1 edition. The book was printed unchanged in many further editions. The ISBN refers to a current reprint of the second edition.

Easwaran, K. (2008). The role of axioms in mathematics. *Erkenntnis*, 68(3):381–391.

Feferman, S., Friedman, H. M., Maddy, P., and Steel, J. R. (1999). Does mathematics need new axioms. *Bulletin of Symbolic Logic*, 6(4):401–446.

Friend, M. (2013). *Pluralism in Mathematics: A New Position in Philosophy of Mathematics*. Dordrecht, Netherland: Springer.

Friend, M. (2019). Varieties of pluralism and objectivity in mathematics. In Centrone, S., Kant, D., and Sarikaya, D., editors, *Reflections on the Foundations of Mathematics: Univalent Foundations, Set Theory and General Thoughts*, pages 345–362. Springer Verlag.

Hamkins, J. D. (2012). The set-theoretic multiverse. *Review of Symbolic Logic*, 5(3):416–449.

Hanson, N. R. (1958a). The logic of discovery. *Journal of Philosophy*, 55(25):1073–1089.

Hanson, N. R. (1958b). *Patterns of Discovery: An Inquiry Into the Conceptual Foundations of Science*. Cambridge University Press.

Hellman, G. (2003). Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 11(2):129–157.

Hellman, G. and Bell, J. L. (2006). Pluralism and the foundations of mathematics.

Koellner, P. (2009). Truth in mathematics: The question of pluralism. In Bueno, O. and Linnebo, , editors, *New Waves in Philosophy of Mathematics*, New Waves in Philosophy. Palgrave Macmillan, London.

Ladyman, J. and Presnell, S. (2016). Does homotopy type theory provide a foundation for mathematics? *British Journal for the Philosophy of Science*, page axw006.

Lange, M. (2014). Aspects of mathematical explanation: Symmetry, unity, and salience. *Philosophical Review*, 123(4):485–531.

Lange, M. (2018). Mathematical explanations that are not proofs. *Erkenntnis*, 83(6):1285–1302.

Linnebo, O. and Pettigrew, R. (2011). Category theory as an autonomous foundation. *Philosophia Mathematica*, 19(3):227–254.

Maddy, P. (2011). *Defending the Axioms: On the Philosophical Foundations of Set Theory*. Oxford, England: Oxford University Press.

Maddy, P. and Meadows, T. (2020). A reconstruction of steel's multiverse project. *Bulletin of Symbolic Logic*, 26(2):118–169.

Maddy, P. and Väänänen, J. (2023). Philosophical uses of categoricity arguments.

Maudlin, T. (2007). *The Metaphysics Within Physics*. Oxford, GB: Oxford University Press.

Mayberry, J. (1994). What is required of a foundation for mathematics? *Philosophia Mathematica*, 2(1):16–35.

Mayberry, J. P. (2000). *The Foundations of Mathematics in the Theory of Sets*. Cambridge University Press.

McSweeney, M. M. (2019). Following logical realism where it leads. *Philosophical Studies*, 176(1):117–139.

Meadows, T. (2013). What can a categoricity theorem tell us? *Review of Symbolic Logic*, (3):524–544.

Moriarty, C. M. (2018). The ad hominem argument of berkeley's analyst. *British Journal for the History of Philosophy*, 26(3):429–451.

Pedersen, N. J. L. L. (2014). Pluralism × 3: Truth, logic, metaphysics. *Erkenntnis*, 79(S2):259–277.

Priest, G. (2001). Logic: One or many. In Brown, B. and Woods, J., editors, *Logical Consequence: Rival Approaches*, pages 153–179. Oxford: Hermes Scientific Publishers.

Priest, G. (2005). *Doubt Truth to Be a Liar*. Oxford, England: Oxford University Press.

Priest, G. (2019). From the foundations of mathematics to mathematical pluralism. In Centrone, S., Kant, D., and Sarikaya, D., editors, *Reflections on the Foundations of Mathematics: Univalent Foundations, Set Theory and General Thoughts*, pages 363–380. Springer Verlag.

Quine, W. V. (1969). Epistemology naturalized. In *Ontological Relativity and Other Essays*. New York: Columbia University Press.

Schlimm, D. (2013). Axioms in mathematical practice. *Philosophia Mathematica*, 21(1):37–92.

Shapiro, S. (2005). Categories, structures, and the frege-hilbert controversy: The status of meta-mathematics. *Philosophia Mathematica*, 13(1):61–77.

Shapiro, S. and Lynch, M. (2019). Does logical pluralism imply, or suggest, truth pluralism, or vice versa? *Synthese*, 198(Suppl 20):4925–4936.

Steel, J. R. (2014). Gödel's program. In Kennedy, J., editor, *Interpreting Gödel*, pages 153–179. Cambridge Univ. Press.

Steinberger, F. (2019). Logical pluralism and logical normativity. *Philosophers' Imprint*, 19.

Steinhart, E. (2002). Why numbers are sets. *Synthese*, 133(3):343–361.

Univalent Foundations Program, T. (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. `https://homotopytypetheory.org/book`, Institute for Advanced Study.

Weyl, H. (1935). Emmy noether. *Scripta Mathematica*, 3:1–20.

Zach, R. (2023). Hilbert's Program. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2023 edition.

# Essay II: Mathematics Needs No (Philosophical) Foundation

## Introduction

In recent years there has been a lot of work arguing for or against one axiomatic theory or another as a *foundation* for mathematics. Most notably, category theory has been put against set theory as a possible foundation for mathematics[1], though there have also been discussions around extending set theory[2]. But does mathematics need a foundation at all[3]? This paper argues that it does not.

The term foundation has not always been used consistently across this literature. There are many ways in which axiom systems might be thought of as a foundation for mathematics, the term as it stands is a cluster concept. To slightly clarify my earlier more provocative claim, in this paper I identify a series of ways in which an axiom system might serve as a foundation for mathematics, and then argue that mathematics has no need of such a foundation. I also clarify a number of weaker senses of "foundation" that I take no issue with.

---

[1]See Univalent Foundations Program (2013) and Corfield (2020) for formal introductions to various versions of homotopy type theory. See Mac Lane (1986), Feferman (1977), Mayberry (1977), Hellman (2003), Awodey (2004), McLarty (2004), Linnebo and Pettigrew (2011) and various works by Ladyman & Presnell including Ladyman and Presnell (2016), Ladyman and Presnell (2019) and Ladyman and Presnell (2020). More recently see several papers in the edited collection Centrone et al. (2019).

[2]Again, see the Centrone et al. (2019) collected volume. See Maddy (1993) and Maddy (2011), the collected essays of: Feferman (2000), Steel (2000), Maddy (2000) and Friedman (2000). See also discussion of the multiverse program in Hamkins (2012) and Maddy and Meadows (2020).

[3]Both Putnam (1967) and more recently Wagner (2019) have taken up a position by this name. Both of them argue that mathematics has no need of a philosophical theory of its foundations. My claim is slightly different, though certainly related in spirit: mathematics does not need its axioms to perform certain philosophical tasks.

In particular, I argue against the need for an ontic or epistemic foundation for mathematics. At least, in any strong or philosophically loaded sense of the term.

Another way to put this in congruence with Essay I: *On Axiom Selection* and Essay III: *Instrumentalism & Axiom Selection* is that I reject certain ontic or epistemic criteria for axiom selection. The ontic or epistemic foundationalist[4] argues that for an axiom to be correct it must satisfy some epistemic or ontic condition. For example, an axiom system must, say, correctly describe the objective mathematical facts, or perhaps an axiomatic system must autonomously justify all of mathematics. By arguing against the need for an ontic or epistemic foundation, this paper argues against the adoption of these criteria as part of a theory of axiom selection.

## Preminimaries

Three preliminaries.

First, I am not rejecting the need or use of foundational work in mathematics, or against an axiomatic system serving as a foundation for mathematics in this purely formal mathematical sense. This is very important and productive work, and not something I would wish to argue against.

In Essay III: *Instrumentalism & Axiom Selection* and Essay IV: *Proofs, Derivations and Axiomatic Pluralism*, I question the need for a *singular* formal foundation for mathematics, advocating for a plural foundation instead. But the merely formal relationship between axioms and ordinary mathematical knowledge is not the subject of this paper, which focuses merely on the (lack of) need for a philosophical foundation for mathematics.

[4]Some terminology: An ontic or epistemic foundation is an axiomatic theory that plays some particular ontic or epistemic role for mathematics. The details of what this role might be are outlined in the relevant sections. Ontic/epistemic foundationalism is the view that mathematics needs an ontic/epistemic foundation. An ontic/epistemic foundationalist is a person who endorses ontic/epistemic foundationalism.

Second, there is a distinction to be made between an axiomatic theory serving as a global foundation and serving as a local foundation. For instance, Peano Arithmetic might serve as a local foundation for arithmetic whereas the set theory-category debate is about which axiomatic theory should serve as a global foundation for the whole of mathematics.

Whilst this paper discusses only global foundations, the arguments made here typically apply to local foundations as well, though perhaps there may be special limiting cases. The claims made should primarily be taken as claims about global foundations, not foundations for specific mathematical sub-disciplines. I welcome these arguments being applied to local cases too, though due diligence should be done to ensure that there aren't exceptional factors that undermine the success of my arguments in those cases.

Last, Penelope Maddy (2000) has pointed out that the term necessary might be unnecessarily strong. Even if mathematics does not need a foundation, it might nevertheless benefit from one. Maddy's point is reasonable, and although I will continue to use the term "need" rather than some weaker normative term, the arguments I present here apply equally well to the claim that mathematics does not benefit from a philosophical foundation as it does to the claim that it does not need one. A reader is welcome to substitute other normative terms in place of "need" as they desire. The argument should remain undisturbed.

## Ontic Foundationalism

The first version of foundationalism to be discussed is ontic foundationalism. The ontic foundationalist holds that axiomatic theories play an ontological role. They entail the existence of any relevant facts about the mathematical objects. In other terms, the ontic foundationalist argues for an ontological criterion on axiom selection. An axiomatic theory is correct only if it correctly describes the mathematical facts[5].

There is a weak and a strong version of this position. I object only to the strong version of ontic foundationalism.

[5]Maddy's Robust Realist is an archetypal example of an ontic foundationalist (Maddy (2007) and Maddy (2011)).

## Weak Ontic Foundationalism

A common distinction in the literature on the role of axioms in mathematics is between what Schlimm (2013) calls the prescriptive and descriptive use of axioms. This follows earlier similar distinctions by Hellman (2003), Shapiro (2005) and others.

Used prescriptively, axioms define a (type of) structure. They determine the kinds of mathematical objects that exist within a certain category. To use the simplest example, the axioms of group theory define what it is to be a group. In contrast, when used descriptively, an axiomatic theory systematizes and describes a fixed body of mathematical facts. The axiomatisation of arithmetic, for instance, can be thought of as a systematization of pre-axiomatic arithmetic.

Note that this distinction is not between axiomatic theories themselves but rather how they are used; the sort of role that one wishes the axioms to play. One might have the philosophical conviction that there really exist mathematical entities called Groups and that Group Theory is the descriptively correct theory of those objects. This would be a descriptive use of a typically prescriptive theory. Similarly, in an alternative history with no pre-axiomatic arithmetic, the axioms of Peano Arithmetic might have been introduced as the definition of a number structure. This would be a prescriptive use of a typically descriptive theory.

There is a sense in which you can think of prescriptive axiomatisation as providing a kind of ontic foundation for their domain. They do, after all, populate a theoretical space with the objects under consideration.

This would, however, be a very weak version of Ontic Foundationalism, and a very minimal constraint on axiom selection[6]. It would not, for instance, actually require that any mathematical objects really exist or that the axioms be true in a non-definitional way.

A nominalist can accept that axiomatic theories define a certain kind of structure, populating said structure with objects, whilst holding that the structures defined are merely possible or perhaps fictional.

[6]The very minimal constraint that axiomatic theories entail the existence of some objects, even if those objects aren't necessarily to be understood as literally real, has been advanced by Hellman (2003) replying to Awodey (1996), who in turn replied in Awodey (2004).

That's not to say that only nominalists can be weak ontic foundationalists. The position is both compatible and in the spirit of a number of realist positions. Classical constructivists like Brouwer (Brouwer (1975), Brouwer (1976), Brouwer (1981) and Brouwer (1992)) or coherentists like Shapiro (1997a) or Rayo (2020) might naturally adopt this kind of a view. The example of coherentism is instructive. The coherentist holds that all logically possible mathematical structures exist. There can then clearly be no axioms (save perhaps the rules of logic) which describe *all* mathematical structures and the task of describing some mathematical structure is rather trivial, just a matter of consistency. Instead, axioms can be thought of as specifying a narrower class of abstract structures are for consideration, a practical necessity given their multitude.

Weak Ontic Foundationalism is a very minimal position. It doesn't require anything philosophically substantial of axiomatic theories. It can be reconciled with almost any view of the nature of mathematical truth. Consequently, this is not a version of ontic foundationalism that I take any issue with.

## Strong Ontic Foundationalism

In contrast to the weak view, a strong ontic foundationalist (henceforth just "ontic foundationalist") holds that an axiomatic theory is correct only if it is a correct description of the mathematical facts or objects. The facts, or at least a portion of them, are determined prior to axiomatization. Axiomatic theories are (in part) descriptive theories of those facts.

A clear analogy here is with scientific theories. Scientific theories, at least according to scientific realists, are descriptive theories of the natural facts. A scientific theory is correct iff it is an accurate description (and perhaps explanation) of the natural facts. For the ontic foundationalist, the relation between scientific theories and natural facts is equivalent, or at least comparable to, the relationship between axiomatic theories and the mathematical facts. An axiomatic theory is correct, or at least correct as a foundation, only if (perhaps iff) it is an accurate description of the mathematical facts.

The first incompleteness theorem is a limiting factor here. No finite, consistent, sufficiently powerful, first-order mathematical theory can describe *all* the mathematical facts. The ontic foundationalist might hold that this means that the correct axiomatic theory is therefore infinite and that all we can know are partial fragments of this theory. They might instead claim that the correct axiomatic theory is second order. Or they might weaken the above requirement such that an axiomatic theory is correct only if (perhaps iff) it is an accurate partial description of the mathematical facts. Which option they might take isn't relevant for this paper.

The clearest recent example of Ontic Foundationalism is Mayberry (Mayberry (1994) and Mayberry (2000)) It is also discussed and rejected by Maddy: her "Robust" Realist is a very clean example of Ontic Foundationalism (Maddy (2007) and Maddy (2011)), and discussed agnostically by Ladyman and Presnell (2016). Historically Frege (Frege (1879) and Frege (1884)) and Gödel (Gödel (1947) and Gödel (1953)) are clear Ontic Foundationalists. Ontic Foundationalism is also frequently tacitly assumed in a number of works by realist philosophers of mathematics - Resnik (1997), Shapiro (1997b), Parsons (2007) and Linnebo (2018), to list just a few examples. Linnebo is most explicit about this, devoting an entire chapter of *Thin Objects* to a demonstration that his iterated abstraction principles produce a model of *ZFC*.

Two objections to ontic foundationalism.

**Objection 1: The Assumption of Realism**

Ontic Foundationalism entails either mathematical realism or axiomatic nihilism (the view that there are no correct axiomatic theories). If ontic foundationalism is true, then there is a correct axiomatic theory iff there are at least some objective mathematical facts that are correctly described by said theory.

I assume that axiomatic nihilism is an untenable consequence of any view. It would always be better to change one's theory of axiom selection than accept axiomatic nihilism. What "correct" might mean will vary from view to view[7] but the answer to the mathematician's question *What axioms should I use?* cannot be *None*.

[7] In particular, the term "correct" does carry some realist connotations. One might wonder if a nominalist can talk of correct axioms. But there is a normative sense of "correct" present in the phrase "the correct decision". This sense is available to the nominalist.

Granting this, I hope, unproblematic assumption, ontic foundationalism entails realism. If there is a correct axiomatic theory iff there are some objective mathematical facts described by said theory and there is, in fact, at least one correct axiomatic theory (Nihilism is false), then there are some objective mathematical facts. But if there are objective mathematical facts, mathematical realism is true.

This leads to a series of objections to ontic foundationalism.

First of all, ordinary nominalist[8] objections to mathematical realism transfer to ontic foundationalism.

If ontic nominalism is true, then there are no mathematical objects. But ontic foundationalism entails that there are (modulo the rejection of axiomatic nihilism) so if ontic nominalism is true, ontic foundationalism is unsurprisingly false.

The epistemic nominalist holds that if realism is true, there can be no mathematical knowledge. But if there is an ontic foundationalism is true, then we couldn't know what the correct axiomatic theory is as we wouldn't be able to know if it correctly describes the mathematical facts or not. But we do have knowledge of which axioms are correct, so ontic foundationalism is false.

That nominalists of various stripes reject ontic foundationalism is not all that surprising. The case for nominalism is also not uncontroversial and the arguments for and objections against nominalism have been well rehearsed, and there's no need for me to repeat them here.

As a member of the quarter of philosophers of mathematics that self-identify as nominalists I, unsurprisingly, take these objections to be a serious issue for ontic foundationalism. However, it would be unfortunate to tie the entire case against ontic foundationalism up in a debate that shows no signs of being resolved any time soon. Consequently, some more theoretically neutral arguments against ontic foundationalism should be found.

[8]In some North American circles the term "nominalism" denotes a specific position outlined by Field (1980). I use the term in line with Burgess and Rosen (1997)'s work. An ontic nominalist rejects the existence of abstract objects. An epistemic nominalist rejects the possibility of knowledge of abstract objects. Any North American readers are welcome to mentally substitute the term "mathematical anti-realist" for "ontic nominalist" and "mathematical sceptic" for "epistemic nominalist".

Fortunately, there are ways to take and generalize the nominalist's objections such that they don't rely on the controversial assumptions of ontic or epistemic nominalism. There are ways to take these arguments and put them on some kind of common ground or, at least, less controversial ground. The argument runs as follows.

Mathematicians should not have to be concerned with philosophical questions regarding realism. Mathematicians should be able to remain agnostic on philosophical questions and this shouldn't disrupt their work.

Easwaran (2008) has argued that one of the functions of axiomatization is to shield mathematicians from philosophical questions. Mathematicians need not answer whether or not their axioms are really true, or the objects described by their theories really exist. They can simply point to their axioms as a particle starting point and go from there.

Similar sentiments are described in Kenneth Kunen's discussion of the philosophy of mathematics in his well-known set theory textbooks. In his early 1980 book, Kunen (1980) adopts platonism (Realism, in my terminology) as the "official" philosophy of the book. By the 2011 version (Kunen, 2011), he had switched to formalism. In both cases, though, Kunen advocates for these positions on *pedagogical* grounds. These are simply the philosophical views that he takes to be most effective in explaining certain crucial points to learners of Set Theory. He's keen to specify in both versions of the book that one can understand what's going on in Set Theory just as well from both of these (and other) perspectives. Mathematical practice is independent of these philosophical debates and mathematicians don't need to take a philosophical stance in order to do mathematics.

But this leads to the following argument against ontic foundationalism. If ontic foundationalism is true and there's an ontic condition on axiom selection, then mathematicians need to make substantive philosophical commitments when adopting a particular axiomatic theory. They need to accept that the axioms are objectively true, and the objects entailed by them or contained within them in fact really exist. But mathematicians do not need to make such commitments, pace practical agnosticism, so there can be no ontic condition on axiom selection and ontic foundationalism is false.

Note the normative requirement here is in fact relatively light. One does not need to make the claim that mathematicians must or even should remain philosophically neutral. It is sufficient for this argument that mathematicians may remain agnostic on metaphysical or philosophical questions whilst nevertheless adopting one axiom system or another for their practical purposes.

If it's even the case that mathematicians simply *may* remain neutral, then ontic foundationalism is false as it forces mathematicians to take a philosophical stance during axiom selection where one is not needed.

**Objection 2: Historical Examples of Axiom Selection**

Ontic Foundationalism cannot account for how axiom selection actually progressed in a number of key historical cases. If ontic foundationalism is true, then mathematicians have behaved in a very puzzling manner at a number of points in history.

To show this, consider the comparison with natural science more closely. Scientific theories are descriptive theories of the natural world. Scientific theories are correct because they describe and perhaps explain the natural facts. To justify their theories, therefore, scientists must go out and collect data, and perform observations and experiments. According to a simple picture, scientists make an observation, use these observations to justify some general claim, and then defend a theory on the grounds that it best explains these observations.

This picture is overly simplistic and the work on scientific explanation over the last half century has shown that a more complex picture is required in practice. For example, what an experiment is taken to signify is not theory-neutral. This is the Duhem-Quine thesis (Duhem (1998) and Quine (1998)). Evidence is not theory-neutral. Moreover, theory guides experimentation by suggesting experiments to run. It is not simply an after-the-fact systematization of the natural facts and regularities that have been discovered. Lastly, the exact methods of science differ greatly from one domain to another. The "simple picture" outlined above is true only at some level of abstraction when looking across the range of methodological diversity within science.

That notwithstanding, an essential component in the scientific method is the abduction from first-order natural facts established on the basement of experimentation to some kind of general theory which explains those facts. None of the issues outlined above disagree with that, they just point out that it's not the entire picture.

If ontic foundationalism is true then, as discussed above, a strong analogy between scientific theories and axiomatic theories holds. Axiomatic theories relate to mathematical facts in much the way scientific theories relate to natural facts. If ontic foundationalism is true, one might expect historical cases of axiom selection to follow this "simple picture" of the scientific method, at least at some appropriate level of abstraction.

Some historical cases do seem to work like this. Both geometry and arithmetic existed in some kind of pre-axiomatic form. The pre-axiomatic geometric or arithmetic facts serve as the analogue to the experimentally identified natural fact. An axiomatic theory is introduced as a way of systematizing, organizing, and explaining these facts. Now, assuming there's some kind of external justification for these pre-axiomatic facts, for instance via their utility in science, then the analogy holds very well.

However, these kinds of cases are of an exception and this does not generalize how axiom selection has proceeded more generally, and certainly not how axiom selection has functioned within the last century or so.

The most detailed account of axiom selection in mathematical practice is provided by Penelope Maddy. According to Maddy, in practice, axioms are adopted because they advance certain mathematical goals. For instance, the axiom of choice was adopted because it had a number of productive and helpful consequences in a wide range of domains and mathematics, advancing a number of mathematical goals.

Maddy's account generalizes to a number of cases. She discusses the current case for large cardinal axioms at length (both Maddy (2007) §4.3 and Maddy (2011)), but given that I take this to be an ongoing rather than resolved case, it's best to look for historic examples for present purposes.

The adoption of Peano Arithmetic can be easily understood on Maddy's account as succeeding in advancing the goal of formalizing arithmetic. The adoption of the axioms of group theory can be understood via the goal of developing an exceptionally useful algebraic tool with a wide range of applications. The rejection of the parallel postulate in geometry can be understood via the goal of moving geometry from a synthetic to an abstract discipline.

In Essay III: *Instrumentalism & Axiom Selection*, I question Maddy's account on a number of grounds. I offer an alternative account in its place where axioms are adopted based on the utility of the structures they produce. But this is a similar enough account to Maddy's and hers serves sufficiently for present purposes.

Assuming Maddy's account of these cases is (more or less) correct, can the Ontic Foundationalist explain the goal-directed method, if mathematical methods should mirror the scientific method?

There is still some level of similarity between the goal-directed method and the simple story of the adoption of scientific theories told above. Novel consequences of particular axioms or axiomatic theories are identified. These consequences are then evaluated on the basis of how well they advance certain mathematical goals. If a critical mass of interesting goal-advancing consequences are identified, then the axioms are adopted.

The explanation holds if evaluating the consequences of an axiom on the basis of its ability to advance certain mathematical goals can be seen as sufficiently analogous to evaluating a scientific theory on the basis of its ability to describe, predict and explain natural facts.

But this is, I think, a difficult comparison to take seriously. It's unclear why we should expect reality to be cooperative with our goals (or those of mathematicians). Why should we expect objective mathematical facts to correspond nicely to the goals of mathematicians?

It would not be acceptable for natural scientists to adopt theories on the basis that they advance certain internal scientific goals, for instance, certain aesthetic considerations. String Theory, for instance, has been heavily criticized on exactly these grounds.

But if ontic foundationalism is true and the analogy between axiomatic theories and scientific theories holds, then one should react similarly to the use of internal mathematical goals in deciding between axiomatic theories.

The ontic foundationalist must then do one of two things in reply to this objection.

First, they could argue that mathematicians have been mistaken in the kinds of methods they have applied. I don't take this to be an entirely unconscionable move, though an exceptionally strong case for ontic foundationalism should be made if it is to be this revisionary of mathematical practice. No such case is forthcoming.

Second, they could explain why, in fact, the historical examples outlined are contiguous with the scientific method, despite clear prima facie differences. There are some potential ways to do this. Maddy's discussion of thin realism in her book Second Philosophy (§4.4) aims to explain how realism is consistent with her acceptance of the methods. She draws a distinction between robust and thin realism, arguing that only robust realism faces the problem outlined here.

I'm sceptical of the success of Maddy's move, though won't critique it here. Suffice, for present purposes, to say that Maddy's thin realist very clearly rejects ontic foundationalism. Realism, for the thin realist, is an uninteresting byproduct of the methods of mathematics. That there are mathematical objects is true simply because the properly selected axioms entail that there are objects. But not then that the thin realist, at least of Maddy's sort, is not an ontic foundationalist. Axiom selection for Maddy has nothing to do with realism, contra the ontic foundationalist's claims. None of this is an unwelcome consequence for Maddy, who is quite clear in her adoption of a goal-based account of axiom selection.

Maddy's thin realism, at best, reconciles the actual methods of axiom selection with realism. It does not reconcile the actual methods of axiom selection with ontic foundationalism. However, nothing I've said here entirely rules out the possibility of a different kind of thin realism that can take on the stronger burden of reconciling ontic foundationalism with the historical cases outlined. Such a view is neither obvious nor forthcoming, I am certainly sceptical that it is possible, but I welcome ontic foundationalist attempts to prove me wrong!

## Epistemic Foundationalism

A simple story of the epistemology of mathematics goes like this. Ordinary mathematical statements are justified by proofs. Proofs are, in essence, short-hand descriptions of derivations from previously established theorems and lemmas, or from the axioms themselves. One could take a proof of a theorem, take the transitive closure of its citations (modulo a little tidying up), unpack the short-hand of the resulting script and in order to reach a deductively valid derivation from the axioms to the theorem.

Work by, amongst others, Azzouni (2004), Burgess and Toffoli (2022), Hamami (2022) and Tanswell (2015) has questioned that such a direct link between proofs and derivations exists. I discuss this further in Essay IV: *Proofs, Derivations and Axiomatic Pluralism*, with particular emphasis on the relationship between proofs, derivations and the axioms.

But even granting a more complex relationship between proofs and derivations, it's still true that whatever justification proofs provide for ordinary mathematical claims, this is transmitted from the axioms.

For proofs to justify ordinary mathematical statements, the axioms must themselves be justified in order to transmit that warrant via proofs[9].

Epistemic foundationalism is the claim that mathematics needs an axiomatic theory that plays this role - An autonomously or independently justified set of statements from which one can transmit warrant via proof to the rest of mathematics.

[9]The nominalist, I note, does not deny this account of the epistemology of ordinary mathematics. However, they think that mathematical facts have a kind of conditional structure. To say that a mathematical statement $\phi$ is true is just to say that $\phi$ follows from some relevantly selected axioms, or is true in all possible structures modelling these axioms, or is true in the fictional universe described by these axioms, etc. Thus the nominalist accepts the account of ordinary mathematical epistemology but without the need for a further justification of the axioms themselves.

## Weak Epistemic Foundationalism

As with ontic foundations, there are weaker versions of epistemic foundationalism that I don't wish to criticize.

First of all, the mere view that proofs need to be tracked back to the axioms could, in a sense, be called a kind of epistemic foundationalism. Some axioms are accepted by the mathematical community as being the basis of proofs. In this practical sense, these axioms do serve as a kind of epistemic foundation. I call this the formal sense of foundation.

I have no criticism of axioms playing this kind of minimal foundational role. This is, again, a very minimal sense of foundation, if it could even be called a foundation at all.

Second, a weak notion of an epistemic foundation that I take no issue with is the idea that axiomatic theories contribute in some way to mathematical knowledge. Dirk Schlimm (2013) has argued, for instance, that we shouldn't think of axiomatic theory as merely an after-the-fact systematization of an already established body of knowledge, but rather axiomatic theories play a role in generating new mathematical discoveries.

I agree with Schlimm's analysis here. This is comparable to the way, for instance, scientific theories might play a role in scientific knowledge generation: identifying new research questions, clarifying important distinctions, etc.

I'm unsure if this should be called a foundation at all. I would describe this more as a *contributory* epistemic role rather than a *foundational* one. However, if this can be thought of as a kind of foundational role, it's a weak notion of foundation that I do not object to.

## Strong Epistemic Foundationalism

As mentioned above, strong epistemic foundationalism (henceforth just epistemic foundationalism) holds that mathematics needs an axiomatic theory that is, in a certain sense, independently justified that can transmit its justification to the rest of mathematics, or at least to the domain that it serves as a foundation for[10].

The most extreme example of epistemic foundationalism is the classical claim that axioms are self-evident. I'll note that there is no one that really holds this view in practice now. Self-evidency is an exceptionally high bar and something that I don't think any axioms in practice would be able to meet.

I take self-evidency to mean not merely a priori demonstrability, but a very immediate kind of a priori demonstrability. Self-evident statements should be justified simply in virtue of understanding their content. There is no gap between comprehension and justification.

This is simply too epistemically demanding. As an example, the Fregian tradition in the philosophy of maths purports to offer a priori demonstrations of the existence of structures modelling certain axioms (see Linnebo (2018) for a recent example and Pearce (2022) for discussion). Even if they're correct, they would still fail to meet the requirements of self-evidency, as these a priori proofs are very complex and certainly not immediate upon understanding the content of the claims. If the most optimistic appraisals of the kind of epistemic access we have to mathematical facts still fall short of self-evidency, then it's clearly too demanding a criterion. A useful example of epistemic foundationalism, but not to be taken seriously.

A more reasonable version of an epistemic condition on axiom selection is Linnebo and Pettigrew (2011)'s justificatory autonomy condition. They begin with a notion of relative autonomy. An axiomatic theory is autonomous relative to another branch of mathematics iff it can be justified without relying on the warrant transmitted from that other branch. An axiomatic theory is autonomous simpliciter iff it is autonomous relative to all branches of mathematics. Mayberry (1994) and Feferman (Feferman (1999) and Feferman (2000)) have advocated for similar positions, though in slightly stronger terms than Linnebo and Pettigrew.

[10] One might imagine a pluralist version of epistemic foundationalism with independently justified foundations for different domains of mathematics.

I'll reply to Linnebo and Pettigrew's weaker formulation of the position though all arguments transfer equally well to the stronger version.

This condition fits with the picture outlined above. Ordinary mathematical statements are justified on the basis of proofs (or derivations) from the axioms. Axioms, however, cannot be justified that way on pain of circularity or a regress. A truly autonomous justification of some axioms would provide just such a justification.

Note that there is a lot of flexibility within the span of this view. Ordinary mathematics might have some degree of independent justification in addition to its justification via the axioms. The level of warrant required of an axiomatic theory to serve as an independent foundation might vary from theory to theory. But, at the least, axioms have some kind of independent justification and they transmit that justification to ordinary mathematical statements via proofs.

Two objections to epistemic foundationalism

**Objection 1: The assumption of realism, again.**

As with ontic foundationalism, the epistemic foundationalist must make an assumption of realism, albeit indirectly. Justification is an indication of truth. So in order for mathematical statements or axioms to be justified, there must at the very least be the epistemic possibility of mathematical objects.

As with the above, I assume that axiomatic nihilism is not an acceptable position. Everyone must hold that there are at least some correct axioms.

Whatever the correct axioms are, the epistemic foundationalist holds that they are (autonomously) justified in believing in those axioms. This is automatically inconsistent with epistemic nominalism, which holds that no mathematical statements are justified (at least, not in the epistemic foundationalist's sense).

Some but not all versions of epistemic foundationalism are inconsistent with ontic nominalism.

There are two important justificatory thresholds that a statement might cross based on the strength of evidence in favour of it. First, a statement might be sufficiently justified to the point where it is merely rationally permissible to believe it, but that it's also rationally permissible not to believe it (either to remain agnostic or actively disbelieve). Second, a statement might be justified to the point rationality requires belief in the statement. Disbelief is not rationally possible, if the strength of the evidence is above the second threshold.

That such a distinction exists is not epistemically uncontentious, I'm merely sketching a view one might hold.

If epistemic foundationalism only requires justification of the axioms to the first threshold (mere permissibility), then it is consistent with ontic nominalism. For a statement justified to the point of mere permission, it is consistent to reject the claim whilst simultaneously holding that one would be justified in believing it.

If epistemic foundationalism requires that axioms be justified to the second threshold (rational obligation), then it is inconsistent with ontic nominalism. One cannot (rationally) claim that one is rationally required to believe X but nevertheless not X is true[11]

So epistemic foundationalism is inconsistent at least with epistemic nominalism and, in some versions, with ontic nominalism as well.

Additionally, nominalists of all stripes will hold that the epistemic foundationalist has simply misunderstood the structure of mathematical statements (this was Putnam's objection to epistemic foundationalism).

[11]The sense of "rationality" used here is merely epistemic. There are other senses of rationality: practical, ethics, etc, where it might be the case that it's rational to believe false things. James (1896)'s Ice Climber case is the classic example. But this is non-epistemic rationality and not the subject of discussion here.

Nominalists generally hold that mathematical statements have some kind of conditional structure (see discussion above). If $\phi$ is some proven mathematical theorem, Nominalists will not take $\phi$ itself to be justified. Rather, if $T$ is the axiomatic theory that $\phi$ was proven from, nominalists might hold that $T \vdash \phi$ is what has been justified (formalism), or that and possible $T$ structure satisfies $\phi$ (modal structuralism) or that $T$ fictions all satisfy $\phi$ (fictionalism), or some other comparable position.

It's these conditional statements, not the theorems themselves, which are justified on the nominalist's view. The epistemic foundationalist, then, simply misunderstands what the content of mathematical knowledge actually is.

As with the previous section, I won't re-litigate the realism vs nominalism debate here. If, like me, one is a nominalist, then these are additional grounds to reject epistemic foundationalism. If one is not, the other arguments are for you. Again, this line of objection can be put on more common ground, and in a similar way.

As argued above, it's reasonable to claim that mathematicians should at the very least be able to remain agnostic on philosophical questions about realism or anti-realism. This is inconsistent with epistemic foundationalism, as it is with ontic foundationalism. The value of justification is only ever as an indicator of truth. Justification is never intrinsically valuable, only ever extrinsically valuable because truth is valuable. If justification is relevant to axiom selection, then this can only be because truth is relevant to axiom selection. But this then violates the claim that mathematics should be neutral towards metaphysical issues, so as to allow mathematicians to be agnostic. Hence epistemic foundationalism is also inconsistent with the permissibility of agnosticism.

All this is largely a repetition of the arguments presented in the previous section, so I simply provide the sketch here and refer back to that section for the details.

**Objection 2: The Wrong Direction of Fit**

For the epistemic foundationalist, axiomatic knowledge is (partly[12]) epistemically prior to ordinary mathematical knowledge. Warrant flows from axioms to their consequences and not vice versa.

Note that this is already an epistemically unusual situation. Typically theories are justified in virtue of having accurate consequences. Explananda are epistemically prior to the explanatia. Now mathematics is epistemically unusual in a number of ways, so it is not especially surprising if this turns out to be another. But this difference should, at the least, raise some prima facie suspicion of the account of mathematical epistemology committed to by the epistemic foundationalist.

These suspicions are further borne out again by looking at historical cases of axiom selection.

Across a wide range of cases, axioms have been justified on the basis of their consequences.

In set theory, the Axiom of Choice became widely accepted because of the wide range of helpful consequences it yields (Zorn's Lemma, all vector spaces have a basis, etc). In geometry, the existence of interesting non-euclidean models was sufficient for the parallel postulate not to be adopted. Peano Arithmetic was justified on the grounds that it provided a good systematization of pre-axiomatic arithmetic.

Above I explain Maddy's interpretation of these kinds of cases. Axioms are evaluated based on the extent to which their consequences advance mathematical goals. As mentioned, elsewhere I've offered a comparable instrumentalist reading where axioms are selected on the basis of their utility - If the structures produced are interesting or useful tools for mathematics or science.

[12]As outlined, an epistemic foundationalist might hold that ordinary mathematical knowledge has both an independent non-axiomatic justification (e.g. via utility and science) and an axiomatic justification. The epistemic relationship between axiomatic knowledge and ordinary mathematical knowledge is still from the axioms to the ordinary statement, even if that's just a full picture of the epistemology of mathematics.

There's also a realist reading of these cases. Comparable to the scientific method, mathematicians begin with knowledge of ordinary mathematics. Perhaps via its application to science (Quine, 1981), or perhaps via other means. From this, they're able to infer abductively to an axiomatic theory on the basis of its ability to explain the truth of their established ordinary mathematical knowledge.

Fortunately, I don't need to argue between interpretations and they all agree on a central point relevant to epistemic foundationalism. All of these views agree that the epistemic direction of fit runs from ordinary mathematics to the axioms, contra the epistemic foundationalist's claims. (Some) Ordinary mathematical knowledge is epistemically prior to the axioms. Axioms and axiomatic theories, on this view, are not justified autonomously from the ordinary mathematics they imply.

Note this is not to claim that *all* mathematical knowledge is epistemically prior to that of axioms. An axiom might be adopted on the basis of some of its consequences, but will then entail a number of other results that weren't previously known. In such a situation, an axiom is epistemically prior to some of its consequences but epistemically posterior to others.

Moreover, as discussed above, I agree with Schlimm's analysis that axiomatics often play a productive epistemic role within mathematics.

But at the very least, in these cases, it seems as if axioms of axiomatic theories were justified on the basis of some ordinary mathematical knowledge, not the other way around. This is contrary to the epistemic foundationalist's claims. Epistemic foundationalism, therefore, gets the epistemic direction of fit wrong between axioms and at least some ordinary mathematical knowledge. Epistemic foundationalism, as with ontic foundationalism, fails to give an accurate count of historical cases of axiom selection.

The outlook for epistemic foundationalism is, I think, worse than ontic foundationalism.

For both, I leave open the option of claiming that mathematicians have been wrong in their historical methodology. They are welcome to make this case, though given how radically revisionary this would be, an exceptionally strong case would need to be made. Such a case is not forthcoming.

For the ontic foundationalist, though, I also left open the option of "thin realism" as a way of reconciling the actual methodology of axiom selection, as it exists in mathematical practice, with ontic foundationalism. I'm sceptical about this, but it is a possibility. No such option is available to the epistemic foundationalist, though. Whilst historical examples of axiom selection are difficult to reconcile with ontic foundationalism (I think terminally so), they refute epistemic foundationalism in a far more direct way. The cases are quite literally examples of axioms being justified because of ordinary mathematical knowledge which is directly contra to the epistemic foundationalist's claims.

Whilst I think the outlook for ontic foundationalism is bad, the outlook for epistemic foundationalism is, I think, worse.

## Conclusion

This paper argues against two forms of foundationalism in mathematics: ontic foundationalism and epistemic foundationalism. Mathematics, so I argue, does not need an axiomatic theory to fulfil certain tasks which I take to be merely philosophical, not mathematical, in nature. To put this in other terms, this paper argues against certain epistemic or ontic conditions on axiom selection. The ability to serve as an ontic or epistemic foundation for mathematics should not be used as criteria for axiom selection.

## References


Awodey, S. (1996). Structure in mathematics and logic: A categorical perspective. *Philosophia Mathematica*, 4(3):209–237.

Awodey, S. (2004). An answer to hellman's question: Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 12(1):54–64.

Azzouni, J. (2004). The derivation-indicator view of mathematical practice. *Philosophia Mathematica*, 12(2):81–106.

Brouwer, L. (1975). *Collected works I*. North-Holland, Amsterdam.

Brouwer, L. (1976). *Collected works II*. North-Holland, Amsterdam.

Brouwer, L. (1981). *Brouwer's Cambridge lectures on intuitionism*. Cambridge University Press, Cambridge.

Brouwer, L. (1992). *Intuitionismus*. Wissenschaftsverlag, Mannheim.

Burgess, J. and Rosen, G. (1997). A subject with no object: Strategies for nominalistic interpretation of mathematics. *Philosophical Quarterly*, 50(198):124–126.

Burgess, J. and Toffoli, S. D. (2022). What is mathematical rigor? *Aphex*, 25:1–17.

Centrone, S., Kant, D., and Sarikaya, D. (2019). *Reflections on the Foundations of Mathematics: Univalent Foundations, Set Theory and General Thoughts*. Springer Verlag.

Corfield, D. (2020). *Modal Homotopy Type Theory: The Prospect of a New Logic for Philosophy*. Oxford, England: Oxford University Press.

Duhem, P. (1998). Physical theory and experiment. In Curd, M. and Cover, J., editors, *Philosophy of Science: The Central Issues*, pages 257–279. Norton, New York.

Easwaran, K. (2008). The role of axioms in mathematics. *Erkenntnis*, 68(3):381–391.

Feferman, S. (1977). Categorical foundations and foundations of category theory. In Butts, R. and Hintikka, J., editors, *Logic, Foundations of Mathematics and Computability Theory*, pages 149–169. Reidel, Dordrecht.

Feferman, S. (1999). Does mathematics need new axioms? *American Mathematical Monthly*, 106:99–111.

Feferman, S. (2000). Why the programs for new axioms need to be questioned. In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Field, H. H. (1980). *Science Without Numbers: A Defence of Nominalism*. Princeton University Press, Princeton, NJ, USA.

Frege, G. (1879). *Begriffsschrift: Eine der Arithmetischen Nachgebildete Formelsprache des Reinen Denkens*. Louis Nebert, Halle a.d.S.

Frege, G. (1884). *Grundlagen der Arithmetik: Studienausgabe MIT Dem Text der Centenarausgabe*. Wilhelm Koebner Verlag, Breslau.

Friedman, H. M. (2000). Ordinary mathematics will need new axioms. In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Gödel, K. (1947). What is cantor's continuum problem? In Benacerraf, P. and Putnam, H., editors, *Philosophy of Mathematics: Selected Readings (2nd Edition)*, pages 470–485. Cambridge University Press.

Gödel, K. (1953). Is mathematics syntax of language? In Gödel, K. and Goedel, K., editors, *K. Gödel Collected Works*, pages 334–355. Oxford University Press: Oxford.

Hamami, Y. (2022). Mathematical rigor and proof. *Review of Symbolic Logic*, 15(2):409–449.

Hamkins, J. D. (2012). The set-theoretic multiverse. *Review of Symbolic Logic*, 5(3):416–449.

Hellman, G. (2003). Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 11(2):129–157.

James, W. (1896). ??the will to believe". In Cahn, S. M., editor, *The Will to Believe: And Other Essays in Popular Philosophy*, pages 1–15. New York: Longmans, Green, and Co.

Kunen, K. (1980). *Set Theory: An Introduction to Independence Proofs*. Elsevier, Amsterdam.

Kunen, K. (2011). *Set Theory*. Studies in Logic: Mathematical. College Publications.

Ladyman, J. and Presnell, S. (2016). Does homotopy type theory provide a foundation for mathematics? *British Journal for the Philosophy of Science*, page axw006.

Ladyman, J. and Presnell, S. (2019). Universes and univalence in homotopy type theory. *Review of Symbolic Logic*, 12(3):426–455.

Ladyman, J. and Presnell, S. (2020). The hole argument in homotopy type theory. *Foundations of Physics*, 50(4):319–329.

Linnebo, O. (2018). *Thin Objects: An Abstractionist Account*. Oxford: Oxford University Press.

Linnebo, O. and Pettigrew, R. (2011). Category theory as an autonomous foundation. *Philosophia Mathematica*, 19(3):227–254.

Mac Lane, S. (1986). *Mathematics: Form and Function*. Springer-Verlag, New York.

Maddy, P. (1993). Does v. equal l? *Journal of Symbolic Logic*, 58(1):15–41.

Maddy, P. (2000). Does mathematics need new axioms? In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Maddy, P. (2007). *Second Philosophy: A Naturalistic Method*. New York: Oxford University Press.

Maddy, P. (2011). *Defending the Axioms: On the Philosophical Foundations of Set Theory*. Oxford, England: Oxford University Press.

Maddy, P. and Meadows, T. (2020). A reconstruction of steel's multiverse project. *Bulletin of Symbolic Logic*, 26(2):118–169.

Mayberry, J. (1994). What is required of a foundation for mathematics? *Philosophia Mathematica*, 2(1):16–35.

Mayberry, J. P. (1977). On the consistency problem for set theory: An essay on the cantorian foundations of classical mathematics (i). *British Journal for the Philosophy of Science*, 28:1–34.

Mayberry, J. P. (2000). *The Foundations of Mathematics in the Theory of Sets*. Cambridge University Press.

McLarty, C. (2004). Exploring categorical structuralism. *Philosophia Mathematica*, 12(3):37–53.

Parsons, C. (2007). *Mathematical Thought and its Objects*. Cambridge University Press, New York.

Pearce, G. R. (2022). A nominalist alternative to reference by abstraction. *Theoria*, 1:1–12.

Putnam, H. (1967). Mathematics without foundations. *Journal of Philosophy*, 64(1):5–22.

Quine, W. V. (1981). *Theories and Things*. Cambridge, Mass.: Harvard University Press.

Quine, W. V. (1998). Two dogmas of empiricism. In Curd, M. and Cover, J., editors, *Philosophy of Science: The Central Issues*, pages 280–301. Norton, New York.

Rayo, A. (2020). The ultra-thin conception of objecthood. A version will be published in a special issue of the journal *Inquiry* devoted to Linnebo's book *Thin Objects*.

Resnik, M. D. (1997). *Mathematics as a Science of Patterns*. Oxford University Press, New York ;.

Schlimm, D. (2013). Axioms in mathematical practice. *Philosophia Mathematica*, 21(1):37–92.

Shapiro, S. (1997a). *Philosophy of Mathematics: Structure and Ontology*. Oxford, England: Oxford University Press USA.

Shapiro, S. (1997b). *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press, New York.

Shapiro, S. (2005). Categories, structures, and the frege-hilbert controversy: The status of meta-mathematics. *Philosophia Mathematica*, 13(3):61–77.

Steel, J. R. (2000). Mathematics needs new axioms. In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Tanswell, F. (2015). A problem with the dependence of informal proofs on formal proofs. *Philosophia Mathematica*, 23(3):295–310.

Univalent Foundations Program, T. (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. https://homotopytypetheory.org/book, Institute for Advanced Study.

Wagner, R. (2019). Does mathematics need foundations? In Sarikaya, D., Kant, D., and Centrone, S., editors, *Reflections on the Foundations of Mathematics*. Springer Verlag.

# Essay III: Instrumentalism & Axiom Selection

How should we select new axioms for mathematics?

One way is to think of axiomatic theories as descriptive theories of some fixed body of mathematical facts. Just as scientific theories are correct iff they correctly describe (and perhaps explain) the natural facts, axiomatic theories are correct iff they correctly describe the mathematical facts (assuming there are such things). An alternative is to treat axioms instrumentally. The function of an axiomatic theory is not descriptive but productive. Axiomatic theories give mathematicians a universe to work with, and the function of axioms is simply to provide the most interesting or productive universe possible[1].

The most well-known (arguably) non-descriptivist theory of axiom selection was developed by Penelope Maddy across two books: *Second Philosophy* (Maddy (2007) - **SP**) and *Defending The Axioms* (Maddy (2011) - **DTA**)[2]. On her view axiomatic theories are evaluated on the basis of their ability to advance the goals of mathematicians. Maddy's account is, I think, very effective in application but theoretically problematic. It typically gets the right results, but on grounds that are somewhat questionable. The aim of this paper is to outline a theory of axiom selection that behaves like Maddy's in application but avoids its issues.

[1]This is not to say anything of the ontological status of the objects entailed by axioms. One of the advantages of instrumentalism is that it's neutral with respect to mathematical realism.

[2]Maddy is a philosopher whose views shift comparatively frequently in contrast with other philosophers. This is something commendable though results in the challenge of identifying to which version of her work one ought reply. At least on my reading, Maddy's views on axiom selection are stable across **SP** and **DTA**. In this paper, I'm responding to the Maddy in **SP** and **DTA**. Her more recent work (Maddy, 2019) applies an approach certainly congruent with **SP** and **DTA**. Though there is, I think, I a slightly more epistemic flavour to that work that takes it a little further from my proposal, hence why it's not discussed here.

§1 outlines Maddy's view, introducing two points of critique along the way. §2 raises two further more substantial objections to Maddy's view (the Euthyphro dilemma in §2.1 and my objection to Maddy's Thin Realism in §2.2). My positive proposal, an instrumentalist theory of axiom selection, is introduced in §3 and shown to resolve the issues raised in the previous sections whilst retaining a theory that is extensionally similar to Maddy's. One notable point of divergence between our accounts is in relation to axiomatic pluralism, the view that there are multiple correct axiomatic theories. I am an axiomatic pluralist and Maddy is not. §3.2 discusses this point of divergence.

# 1 Peneolope Maddy's Theory of Axiom Selection

## 1.1 Three Types of Consideration

Maddy's view is best outlined in her own words. In **SP** she says the following:

> *...The Second Philosopher, in her effort to understand the world, may well turn to the pursuit of pure mathematics, and that when she does so, her assessment of proper methods rests on weighing their efficacy towards her mathematical goals.*
> (**SP** - IV.4, p361)

Mathematicians have a number of goals. Many of those goals relate to mathematics. Mathematicians might wish to see a satisfactory formalization of their intuitive mathematical concepts such as "number", "space" or "set". They might notice certain structural similarities across mathematics and might have the goal of describing these commonalities (see the developments of group and category theory). There might be certain physical phenomena that they wish to model. Axiomatic theories are judged on their ability to advance these goals.

In **DTA** Maddy continues with specific reference to set theory:

> *[The second philosopher] takes the proper methods for introducing sets, for adding new axioms to our theory of them, to be methods of the sort we've been rehearsing: sets are legitimately posited as effective means towards various mathematical goals (in analysis, algebra, foundations and elsewhere); axioms are defended by a careful balance of detailed considerations, both intrinsic and extrinsic.*
> *(**DTA** - II, p56)*

Two new characters have appeared in Maddy's account: intrinsic and extrinsic reasons.

Both are discussed in greater detail below but, in brief, an intrinsic reason is a reason to select an axiom based on features of the axiom *itself* whereas extrinsic reasons are reasons based on what the axiom *does*.

An immediate exegetical question presents itself: What is the relationship between in/extrinsic considerations and goal-directed considerations? Are there three types of considerations: goal, intrinsic and extrinsic? Or are intrinsic and extrinsic reasons specific types of goal-directed considerations?

The latter reading is, I think, more natural. Mathematicians might have certain intrinsic goals, e.g. to explicate certain pre-formal concepts, and they might also have extrinsic goals, goals to do certain things with their axioms.

There is a further question regarding what *the* mathematical goals actually are. Mathematics is not a homogeneous community that can be assumed to have unified aims and goals. Intellectual histories too often fall into the trap of whigism, favouring romantic simple narratives about shared intellectual goals spanning generations over the mess of discord and disagreement. All this is to say that Maddy does not provide an account of how the potentially diverging goals of individual mathematicians are to be aggregated into *The* mathematical goals that can then be used to evaluate axioms.

This is, I think, a serious, though likely solvable, issue for Maddy's position. It's an issue my proposal side steps.

Returning to clarifying Maddy's position, what are intrinsic and extrinsic considerations for adopting an axiomatic theory?

### 1.1.1 Intrinsic Considerations

Intrinsic considerations for adopting axioms are conceptual in nature. They are about how axioms relate to pre-formal concepts. Axiomatic theories are often intended as explications of pre-axiomatized mathematical concepts. Historically, there have been many cases of mathematicians adopting the goal of formally explicating a pre-mathematical concept.

Note that it's axiomatic *theories* that explicate pre-axiomatic mathematical concepts, not individual axioms. Taking Peano arithmetic as an example, the axiom "every number has a unique successor" does not itself exhaust the pre-axiomatic concept of "number". Clock arithmetic, for instance, satisfies this axiom but is not an example of (natural) number. Only together do the axioms of Peano arithmetic (arguably) produce a formal explication of the concept of "number".

Individual axioms, however, are defended on the grounds that they are *implicit* in the pre-axiomatic concept. Whatever an explication of the concept of "number" looks like, it should at least entail unique successors, or else it would be too dissimilar from the original concept and fail to be an explication.

Intrinsic considerations also include a weaker sense in which an axiom might be implicit in a pre-axiomatic concept. Maddy gives the example of inaccessible cardinals. Maddy begins by outlining the well-known iterative conception of sets - the idea that sets are built up in layers. One starts with the empty set, builds all the sets one can out of that (just the set of the empty set), then iterates this. Whenever one approaches a limit, one can take a union at that limit and carry on. A crucial part of the iterative conception of sets is the idea that this process is unbounded and may go on forever. But if one denies the existence of an inaccessible, then this is tantamount to putting an upper bound on the process (namely the inaccessible itself). Maddy claims that therefore implicit in the concept of the iterative conception of sets that there are inaccessible.

Now, I do think Maddy's reasoning can be contested here. At the least, this is a much weaker sense of "implicit in the concept of" than in the case of **PA**'s successor axiom. No formal theory that doesn't contain or entail the successor axiom could ever be a successful explication of the notion of "natural number". It would just be too different from the original concept. A theory which doesn't contain or entail large cardinals could still be a successful *explication* of the pre-axiomatic concept of set; though perhaps an explication inferior in other ways. A better way to put it, I think, is that adopting large cardinals is in the spirit of the pre-mathematical concept of a set, even if not directly implicit in it. The extent to which one ought care about that whilst explicating a concept I leave as an open question.

### 1.1.2 Extrinsic Considerations

Extrinsic considerations justify axioms on the basis of what they do. The Axiom of Choice is the paradigmatic case of an axiom justified on extrinsic grounds. Choice faced an uncertain reception upon introduction into the mathematical world (Moore, 1982). Choice was certainly not an essential part of the pre-axiomatic concept of a set, else it would have been accepted. When Choice was accepted, it was a result of the myriad of useful consequences of Choice discovered by mathematicians.

To put this in terms of goals, there were a series of goals set theorists had relating to what they wanted their set theory to do. This included some element of explication of a pre-axiomatic concept but this didn't exhaust the mathematical goals at play. Set theorists wanted set theory to do certain work for them. Choice is adopted because it helps do that work.

Extrinsic considerations, Maddy admits, are diverse. There seems to be more than a little appetite on Maddy's part for future work clarifying what different types of extrinsic reasons might be. She says the following:

*At this point, even with this limited sampling, it should be clear that a number of different kinds of justifications are being collected together as 'extrinsic'. We have some idea of what's intended by 'intrinsic'... but 'extrinsic' is being applied willy-nilly to any compelling justification that isn't clearly intrinsic.*
*(**DTA** - V, p130)*

My positive account provides just such a unifying story. What are some examples of different kinds of extrinsic considerations?

A recurring motif in Maddy's examples of extrinsic considerations is depth. It's important that new axioms reveal deep connections between mathematical structures. A general type of mathematical goal, according to Maddy, is finding instances of and exploring mathematical depth. An example given later in the book is the connection between $AD^{L(\mathbb{R})}$ and the existence of Woodin Cardinals. These are two different statements independent of $ZFC$ that appear at first sight to have no connection to one another. It turns out that they are, in fact, deeply interconnected. This novel connection is an example of depth that speaks in favour of $AD^{L(\mathbb{R})}$, according to Maddy.

Second, axioms might resolve statements independent of the system they're added to. For instance, some axiom of set theory might resolve the continuum hypothesis. There are two possible sub-variants of this kind of reason.

First, Maddy is keen for set theory as a whole to decide as many statements as possible. Of course, a formally complete and consistent set theory is impossible, due to Gödelian constraints. But, as far as possible within those constraints, Maddy desires a Set Theory that is as decisive as possible. So, on Maddy's view, resolving independent statements seems to be a virtue of an axiom, tout court.

But in other cases, it matters not just that an independent statement is resolved but that it is resolved in a desirable manner. For example, Maddy considers a number of reasons against $V = L$ (See both the discussion in **SP** (§IV.3, p234-261) and earlier in Maddy (1993)). Now, $V = L$ resolves a great many indeterminate statements. One gets the generalized continuum hypothesis (*GCH*) and the axiom of choice. It gives a very neat and well-behaved model of set theory. It was for these reasons that Quine liked it. However, amongst other reasons given,(1) $V = L$ resolves *GCH* the 'wrong' way (many set theorists including Gödel believe it to be false, despite its independence from *ZFC*) and (2) $V = L$ provides an ad hoc restriction at its conception on the possible subsets of $\omega$, considering only the definable subsets.

Note also that it's not just that $V = L$ is ad hoc, it's that the structures it rules are interesting in certain ways. The structures are interesting in a certain way, they instantiate certain mathematical virtues

There's an interesting relationship here between resolving indeterminate statements and providing interesting mathematical structures. Resolving any indeterminate statement is a double-edged sword. Any previously indeterminate statement that is resolved by a new axiom will throw certain models out of consideration, allowing mathematicians to focus on what remains. This can be a good thing if the models thrown out misbehave in a certain way, or were in some other sense faulty or 'non-standard'. Choice, for instance, throws out such models. The resulting theory allows one to say more about the structures one's left with as one doesn't have to contend with aberrant counter-examples (e.g. baseless vector spaces or uncountable unions of countably many countable sets). But, on the other hand, models thrown out might contain interesting structures worth exploring. So-say aberrant counter-examples that are thrown out might, in fact, not be aberrant but have interesting properties. $V = L$ is too restrictive; it throws out too much.

There's then a balance to be found. Selecting one's axioms involves a careful process of inclusion and exclusion. One wishes to exclude as many aberrant models as possible, whilst not discarding interesting structures.

An analogy I find helpful here is with gardening. When pruning a shrubbery one needs to be careful. With insufficient pruning, the shrubbery becomes overgrown with undesirable weeds or overgrown plants but over-cutting risks discarding healthy, beautiful plant life. The good horticulturalist knows how to cut back just the right amount; to trim weeds or overgrown plants but leave healthy shrubs and flowers intact.

The development of non-euclidean geometry is a nice and well-known example of this point. When faced with the independence of the parallel postulate, geometers had two choices. They could adopt the parallel postulate as a new axiom and explicitly narrow the sphere of consideration to Euclidean models or they could reject the parallel postulate and contend with the wide range of possible models that this allowed. They did the latter because the non-euclidean models were incredibly rich and interesting with all sorts of utility inside and outside of mathematics.

All of the above are examples of justifying axioms on the basis of their consequences *within* mathematics. There are also extrinsic considerations that motivate axioms based on their utility outside of mathematics.

A paradigm case, as before, is non-euclidean geometry. One of the main differences between Newtonian and Einsteinian physics is the shift from treating space as Euclidean to non-euclidean. If the geometer's choice to exclude the parallel postulate on the grounds that non-euclidean geometries are mathematically interesting was insufficient, then the fact that they are physically interesting as well should add further weight to the case.

That being said, Maddy's very clear that she puts much more weight on mathematical reasons (both intrinsic and extrinsic) than on reasons of scientific utility. This is where she breaks from Quine (1981). But she's also clear that utility for science is an important goal of mathematics, and in fact the very reason her second philosopher becomes interested in mathematics in the first place. This is certainly a type of reason for Maddy, but not an especially strong one.

Another question is if any open debates in axiomatics could actually be resolved via considerations of utility for science. Systems much weaker than *ZFC* are already enough for modern science. Perhaps there's some possible result in Physics that requires or can be better explained given some number with specific properties inconsistent with *CH*, but I'm unaware of one if there is.

#### 1.1.3 In Summary...

Maddy claims the following about axiom selection:

1. Axioms or axiomatic theories are adopted on the grounds that they advance certain mathematical goals.
2. These can be divided into intrinsic and extrinsic grounds.
3. Intrinsic: The axiom is implicit in some pre-axiomatic mathematical concepts.
   (a) This type of goal is most relevant when mathematicians are involved in a project of explication of a pre-axiomatic concept.
   (b) Sometimes the axiom is *essential* to the reasonable explication of the concept (e.g. the successor axiom in PA).
   (c) Sometimes the axiom is inessential but *in the spirit* of the concept to be explained (e.g. inaccessible cardinals)
4. Extrinsic: The axiom has helpful *consequences* which help to advance the various mathematical goals.
   (a) Mathematical: This is a consequence inside mathematics. Typically this means entailing the existence of mathematically interesting structures (whilst not eliminating structures of equal or greater interest). (E.g. the axiom of choice)
   (b) Non-Mathematical: This is a consequence outside mathematics, presumably in science (e.g. the rejection of the parallel postulate, though there were also mathematical reasons for this).

## 1.2 Intrinsic vs extrinsic reasons

How should intrinsic and extrinsic reasons be balanced against one another? At the very end of **DTA**, quite literally in the final subsection, Maddy argues against Feferman's view (Feferman (1999) and Feferman (2000)) that intrinsic justifications are either (1) the only proper justifications for axioms or (2) better justifications than extrinsic. Maddy argues for quite the opposite. She argues that really its extrinsic justifications (though she might be including goals here) provide the fundamental normative force for axiom selection. Intrinsic justifications are only instrumentally valuable towards extrinsic reasons.

She provides two reasons for this.

One reason is that having a clear mathematical concept gives one evidence that the axioms formalizing that concept are consistent. Consistency is an important mathematical goal[3]. However, this virtue of intrinsic reasons can't explain the value of intrinsic reasons in a number of apparently good cases. The inference from the iterative conception of sets to small large cardinals can't easily be understood in this way. *ZFC* is of strictly weaker consistency strength than *ZFC* plus an inaccessible. Adding inaccessibles decreases one's confidence in consistency.

The second, and I think more important, explanation of the priority of extrinsic reasons is their role in shaping the development of mathematical concepts.

Concepts, in general, can change. Mathematical concepts are no different and there are many instances of mathematical concepts forming and changing. The most famous example is from the development of set theory itself. Amongst the many changes in mathematics that occurred during that period, the concept of a set shifted. Fregean set theory was built on a naive conception of a set, that sets were just extensions of concepts. But the discovery of Russell's paradox lead not just to the rejection of Fregean Set Theory, but to the rejection of the naive conception of sets, in favour of the iterative conception. Another example is, again, geometry. The discovery of non-euclidean geometries lead to a change in the conception of what a geometric space *is*.

[3]The value of consistency can also be understood in extrinsic terms. Inconsistent theories are not mathematically interesting, as they are trivial

Going further, given that mathematical concepts can change, intrinsic considerations look suspect when in conflict with extrinsic considerations. If an intrinsic reason speaks for or against an axiom, this could just as easily be reason to adopt or reject the axiom, rather than reason to simply change the mathematical concept in question. Intrinsic reasons can therefore only support an axiom if there are grounds to think that the underlying concept, from which the reason is derived, is a good one. To justify the concepts, one needs extrinsic reasons[4].

Now, of course, intrinsic considerations have often been historically important. In particular, they are important in cases where a well-developed but pre-axiomatic mathematical concept is formalized. In these cases, it's already known that something like the pre-explicated concept will be incredibly extrinsically useful. It's worth in this context trying to explicate the pre-axiomatic concept in order to have an axiomatic theory that (1) has all the benefits of being a formal and precisely stated theory whilst (2) retains the interesting consequences of the original concept. But then the intrinsic considerations are valuable only instrumentally towards the extrinsic virtues of the original concept[5].

But if, upon axiomatization, one finds that there are formal options that (1) are less faithful to the original concept but (2) nevertheless lead to more interesting structures, the extrinsic reasons have priority.

[4]This is one instance of my general objection to conceptual explication as a philosophical method. See a parallel discussion in relation to logic in Essay IV: *Three Approaches To the Philosophy of Logic*

[5]A short historical tangent. Historically, there seems to have been two ways in which axiomatic theories have developed. On one hand, axiomatizations of Arithmetic and Geometry were attempts at explicating very well developed as pre-axiomatic concepts. This is strange to note in relation to geometry, given that Euclid's *Elements* was published circa 300BC, but this is simply testament to hold old geometry is. Euclid was still building on 3 millennia of pre-axiomatic geometry. On the other hand areas like Group Theory or Homotopy Type Theory developed as formal tools with an intended use in mind. They were pure creations, not explications. Clearly, intrinsic considerations are applicable in the first type of axiomatization but *not* in the second. This distinction I think also loosely tracks Schlimm (2013)'s distinction between descriptive and prescriptive uses of axioms, which I discuss in Essay I, *On Axiom Selection*. Set Theory is interesting as a case because it sits somewhere between the two. There was so pre-axiomatic Set Theory, i.e. uses of an intuitive notion of set prior to the explicit axiomatization of the concept. But there was not nearly so much pre-axiomatic Set Theory as there was pre-axiomatic mathematics or geometry. The pre-axiomatic concept of a set is not precise enough to answer many questions about the foundations of Set Theory. It doesn't appear to resolve choice, or the continuum hypothesis, for instance. The Axiomatization of Set Theory is then torn between these two competing programs: (1) The descriptive goal of explicating the pre-axiomatic concept of a set and (2) the prescriptive goal of creating a new mathematical concept that does certain kinds of foundational work.

Maddy says the following:

> *What's striking is that all these perfectly reasonable ways of proceeding are in fact grounded in their promise of leading to the realization of more of our mathematical goals, to the discovery of more fruitful concepts and theories, to the production of more deep mathematics. Ultimately we aim for consistent theories, for effective ways of organizing and extending our mathematical thinking, for useful heuristics for generating productive new hypotheses, and so on; intrinsic considerations are valuable, but only insofar as they correlate with these extrinsic payoffs. This suggests that the importance of intrinsic considerations is merely instrumental, that the fundamental justificatory force is all extrinsic. This casts serious doubt on the common opinion that intrinsic justifications are the grand aristocracy and extrinsic justifications the poor cousins. The truth may well be the reverse!* (**DTA** - V, p136)

So whilst §1.1 represents Maddy's views as outlined in **SP** and Chapters I-IV of **DTA**, Chapter V of **DTA** gives a subtly different account which puts primacy on extrinsic considerations. This is prima facie difficult to reconcile with her previous claims about axiom selection being a matter of advancing mathematical goals. Mathematicians can, after all, have the kinds of explicatory goals that give rise to intrinsic considerations. What Maddy has to claim is that, as a matter of contingent fact, explication is only ever done for the sake of extrinsic payoff, never for its own sake.

But one wonders what role mathematicians and their goals are now playing at this point in Maddy's picture. If extrinsic payoff in the form of the discovery of some interesting mathematical structure is what's ultimately of interesting, why not simply say that whatever makes these structures interesting (what I later call mathematical virtues) are the goods of axiom selection?

Tying the good of mathematics to the goals of mathematicians embroils one's account in difficult questions about aggregating goals and desires (outlined above) or in the Euthyphro dilemma (§2.1). An account that, so say, "cuts to the chase" therefore seems preferable in virtue of its comparative elegance and simplicity.

Of course, on charge of being ad hoc, a story as to why these particular features of mathematical structures are good is owed. My positive proposal aims to provide just such an account.

## 2 Issues with Maddy's view

Maddy's view does a lot right. In particular, her analyses of particular cases are frequently thorough and insightful. A view that ends up looking, more or less, like hers, especially in how it deals with a lot of these cases, is likely going to be a good one.

My issue is more with the groundwork, the theory behind Maddy's view. My aim is to build a view that's extensionally similar to Maddy's view but avoids issues with the background theory.

Two objections are raised above: (1) Maddy has not provided us with an account of how the goals of individual mathematicians across different time periods aggregate into the singular, unified mathematical goals she needs for her theory (2) Given that a particular goal of finding and understanding interesting mathematical structures seems to be fundamental, why not cut out the middle man and simply claim that whatever makes these structures interesting (what is below called "mathematical virtues") is what's mathematically valuable?

I'll raise two further issues in this section: The Mathematical Euthyphro Dilemma and my issues with Maddy's Thin Realism.

### 2.1 The Mathematical Euthyphro Dilemma

A common type of theory across all normative domains is what I'll call the "ideal agent" approach to normativity. On this approach, one defines the good of that domain in virtue of the aims and goals of certain ideal agents. In the ethical domain, the Good might be understood as that which the virtuous value or aim for (or perhaps that which God values or aims for). Similarly in the ettiquetical domain, one might understand the good of etiquette as that which an ideally well-mannered agent would aim for. Similarly, one might understand the good of rationality as that which rational agents aim for or that which a rational agent would do. The common feature in all of these cases is that the fundamental normative facts are the facts about the goals and actions of an ideal agent. From these facts, all other normative facts are derived. This is a well-known approach to normativity that dates back at least to Aristotle[6].

[6] I avoid using the term "virtue" or "virtue theory" here as it's ambiguous. Virtue theories can be read as an ideal agent theory but it's not the only way. One can also read Aristotle as being, in

All ideal agent approaches to normativity run into variants of the Euthyphro Dilemma. If a particular X is good within a domain because the idea agents value it, then the theory is arbitrary. There is no reason why this particular contingent set of goals that the ideal agent has are the goals that they either must or should have. The ideal agent theorist has one of two options. First, bite the bullet and accept this consequence. Second, concede that facts about the ideal agent are, in fact, not normatively fundamental and there are some other normative facts that explain why the ideal agent's goals are the right ones.

Additionally, there's nothing in this account of ideal agent approaches to normativity that hangs on the agent being singular. The "agent" could be a group of agents instead. One might imagine a polytheistic divine command ethics where the Good is the aggregate opinion of some Council of Gods; the same objection would apply. Is the council's aggregate opinion on X the way it is because X is good, or is X good because of the council's aggregate opinion? Taking the former undermines the ideal agent approach by rejecting the normative fundamentality of the agential facts. Taking the latter makes goodness arbitrary.

If one takes goal-based reasons for axiom selection seriously, this is effectively a grouped ideal agent theory of the norms of axiom selection. The (group) agent is the mathematical community as a whole. The good of the normative domain of axiom selection is that which is aimed at by this ideal agent, i.e. by their goals.

But then, by familiar reasoning, the view falls into a Euthyphro variant. If a feature of an axiom is good simply because it accords with the goals of mathematics, then at least this dimension of this goodness is arbitrary. But if if the goals of mathematics are the way they are because those goals are good (i.e. because mathematicians are sensitive to the good-making features of axioms), then accordance with a goal of mathematics is not something that's per se a good of a particular axiom or axiomatic system.

Of course, Maddy is welcome to bite the bullet and accept some degree of abitrarity in her view. Whilst a comparative disadvantage of a theory, arbitrarity is not terminal. A reading of Maddy that wishes to take her claims about the fundamental value of advancing mathematical goals seriously should likely take this horn of the Euthyphro.

---

essence, a eudaimonia consequentialist.

Nevertheless, something of a reply can be found in **SP** that indicates that Maddy might be inclined in the other direction. Maddy considers the worry that if mathematicians, for whatever reason, adopted strange goals, Maddy would be committed to go along with them. Like Barkley, one might imagine them allowing theological considerations to play a role in mathematics (See Moriarty (2022) and Moriarty (2023) for discussion). This is obviously not the Euthyphro dilemma but is parallel. It still considers how Maddy's commitments modally co-vary with the goals of mathematicians.

Maddy says the following:

*There's nothing in this strange tale told so far to determine whether or not the practice of these wayward souls would continue to be called 'mathematics', and of course the word doesn't matter. What is clear is that the new practice, whatever it's called, wouldn't play the same role in [my] investigation of the world as the discipline we call 'mathematics' now plays. Presumably the evolved practice would end up more or less comparable to 'pure astrology' or 'pure theology' and [I] would have no interest beyond the sociological, anthropological, biological, etc.*
(**SP** - IV.4 p350-351)

One is reading between the lines a little, but Maddy seems to say here that it's not that what she's interested in is advancing the goals of mathematicians *per se*, rather she thinks that actual mathematicians, who are not *"wayward souls"*, have identified something valuable to Maddy's inquiries and have pursued that. In the context of Maddy's wider work, presumably that something is mathematical depth or related notions. In places, when pushed, Maddy seems willing to abandon her ideal-agent account in favour of simply looking to maximize the discovery of deep connections, or other interest-making features of mathematical structures. In other words, taking the second horn of the Euthyphro dilemma. This would, though, be a departure from her position as it appears in the bulk of **SP** and **DTA**.

If Maddy wishes to make this move, as mentioned above, we are owed an account of why mathematical depth or related notions are valuable. Why is this something an enquirer should care about?

I'm going to take the second horn of the Euthyphro dilemma here and attempt to meet this challenge head-on, giving an account of why mathematical depth and related notions are valuable, with references to Instrumentalism about mathematics.

## 2.2 A Nominalist Reply to Maddy's Thin Realism

A final issue with Maddy's view is its reliance on Maddy's Thin Realism (to be explained below). There are, I think, reasons to be sceptical of Thin Realism and it is therefore preferable to motivate an account of axiom selection without relying on Thin Realism. In this section, I provide a brief nominalist reply to Maddy's Thin Realism.

Maddy is interested in explaining the reliability of her methods of axiom selection. This is the subject of Chapter IV.4 in **SP** and Chapter III in **DTA**. She wishes to show that the goal-directed methods she advocates for do, in fact, lead to true mathematical beliefs.

Note, especially in relation to the discussion in Essay I: *On Axiom Selection*, that Maddy is adopting a descriptive condition on axiom selection. This is strange given that the account of axiom selection given thus far has taken goal-advancement (a normative criterion) to be the sole criterion of axiom selection. It's unclear why Maddy feels she needs to meet the challenge of showing why her methods are reliable indicators of truth, instead of simply dropping the descriptive criterion on axiom selection.

Maddy is aware that demonstrating the reliability of her methods is not necessary for adopting her account of axiom selection. In both **SP** Chapter IV.4 and **DTA** Chapter IV Maddy outlines a possible "Arealist" position that has a great deal in common with my Nominalist. The Arealist rejects both (1) the need to show that their mathematical methods are reliable; the methods are *useful* and that's enough and (2) Maddy's Thin Realism. The Arealist position, Maddy acknowledges, is a viable one that has full access to her account of axiom selection. She says the following:

*On [the correct methods of axiom selection] the Arealist and [I] will completely agree; the difference only comes in the way the word 'true' is then applied: [I] find in [my methods] good evidence for truth, while the Arealist takes talk of 'truth' as inessential and sticks to the methodological facts unadorned. As far as the practice of set theory is concerned, it is hard to see what's lost on the Arealist's approach.*

(**SP** - IV.4, p384)

So nothing is at stake here regarding the appropriacy of various set-theoretic methods. Both Maddy's Arealist/my Nominalist and Maddy agree (more or less) on the correct methods of axiom selection. They disagree on if these methods are to be justified by adopting a descriptive criterion on axiom selection along with (thin) realism or by rejecting both and providing an alternative justification.

Maddy grants that traditional realism, what she calls "robust realism" will not demonstrate the reliability of her methods. If mathematical objects are external objects, independent of our conceptualization of them or methods relating to them, to be investigated in a manner more akin to the traditional scientific method, then the reliability of Maddy's goal-directed methods cannot be explained. Traditional "Robust" Realism will not do.

Instead, Maddy opts for what she calls Thin Realism. She says the following:

*[sets are taken to have the properties ascribed to them by set theory and to lack the properties set theory and natural science ignore as irrelevant. There is nothing more to be said about them. Such posits are sometimes called 'thin', so let's call this Thin Realism*
(**SP** - IV.4, p369)

Maddy also quotes Steel who says the following:

*To my mind, Realism in set theory is simply the doctrine that there are sets . . . Virtually everything mathematicians say professionally implies there are sets. . . . As a philosophical framework, Realism is right but not all that interesting.*
(FOM posting 15 Jan. 1998. Printed in **SP** - IV.4, p368 with Steel's permission..)

*'there are sets' . . . is not very intriguing. 'There are sets' is, by itself, a pretty weak assertion! Realism asserts that there are sets, and hence . . . that 'there are sets' is true.*
(FOM posting 30 Jan.1998. Printed in **SP** - IV.4, p368 with Steel's permission.)

So all it takes for a mathematical object to be real is that its existence follows from the methods and practices of mathematics. Similarly, all it takes for a mathematical object to have a certain property is that it's ascribed that property by the methods and practices of mathematics.

My nominalist replies that Maddy has passed the buck. Precisely what was to be proved was that the methods and practices of Mathematics are reliable, Thin Realism simply asserts this fact. But to establish Thin Realism one therefore must establish precisely the point under contention, the reliability of these methods.

What evidence does Maddy have that her Thin Realism is true?

Maddy takes different approaches in **SP** and **DTA**.

In **SP** Maddy embraces a certain level of circularity. She knows that Thin Realism is true because her set theoretic methods have shown her that what is true of sets is all and only those things the methods of set theory say about them. This circularity, though, she takes to be acceptable.

In **SP** (IV.4, p370) she draws a comparison with her reply to general philosophical scepticism. When attempting to establish the reliability of her (non-mathematical) scientific methods, Maddy similarly embraces a level of circularity. She justifies the reliability of her scientific methods using her scientific methods. The sceptic, she holds, is asking her to provide a special kind of "philosophical", non-scientific justification for her methods. Maddy grants that this is something she can't do, but simply rejects that it's something she needs to (**SP**, Chapter I.2). She sees no reason to meet the sceptic on their terms.

Maddy holds that my nominalist's line of criticism makes the same mistake as the sceptic[7]. I am asking her to justify her mathematical methods in a special philosophical way. She grants that this is something she can't do but rejects the need to. But this is not an accurate characterization of my nominalist's reply. My nominalist is not asking Maddy to justify her methods to special philosophical standards, but merely to ordinary (non-mathematical) scientific standards using ordinary (non-mathematical) scientific methods. Maddy accepts in her discussions relating to the use of mathematics in science that scientific methods don't justify belief in mathematical objects (**SP** IV.2 and IV.4, and **DTA** III)[8].

Later Maddy seems to acknowledge this in her comparison of the Arealist (who is similar to my nominalist) and Thin Realist positions. Citing a discussion with Rosen (1999) she states that this comes down to a question of if mathematics is what Rosen called an "authoritative":

[7]Maddy has this objection delivered by a Robust Realist rather than a Nominalist, though the objection is the same.

[8]Of course, an option here is for Maddy to revise her claim that the utility of mathematics in science doesn't provide evidence for the existence of mathematical objects. Maddy would then likely be a Robust Realist. She would then face the (likely insurmountable) challenge of explaining the reliability of her methods, given Robust Realism.

*A practice is authoritative if, whenever we have reason to accept a statement given the proximate goal of the practice, we thereby have reasons to believe that it is true.* (Rosen, 1999, p. 471)

But Maddy and my Nominalist agree that the existence of mathematical objects cannot be established on ordinary scientific grounds. Her methods cannot be justified as reliable in an ordinary scientific science. The question is if notions such as "true" or "exists" should be expanded so as to include mathematical statements (read literally as opposed to nominalistically) and mathematical objects.

I'm not opposed to semantic revision tout court. I discuss semantic revision in the context of logic extensively in Essay V: *Language, Truths & Logics* and Essay VI: *Three Approaches to the Philosophy of Logic*. I defend semantically revisionary programs in logic against a series of objections.

But what Maddy needs to do is explain why these semantic revisions are useful. It's clear the sorts of downsides it might incur: the expansion would be ad hoc; truth and existence would become disjunctive concepts. Two clear downsides. Given that Maddy accepts that nothing here is at stake for mathematical practice, it's unclear what benefits this semantic revision would bring to outweigh these negatives.

Maddy, I think, summarizes the disagreement well:

*In our rough-and-ready terms, we might describe this contrast in yet another way: the Arealist holds that mathematics is distinguished from other extra-scientific enterprises by its role as a handmaiden to science; this very handmaiden role prompts our Thin Realist to assume that mathematics is a science, alongside the various natural sciences.* (**SP** - IV.4, p386)

But Maddy acknowledges that we can account for the role of mathematics in science without recourse to realism. But on those grounds, the move from the utility of mathematics for science to mathematics *being* a science (in the relevant sense) is mistaken.

Thin Realism is then, I think, unjustified. Maddy, by her own admission, cannot justify Thin Realism on ordinary scientific grounds but also, I argue, fails to explain why one should accept the purely mathematical grounds as evidence of their own reliability. Her case comes down to arguing that mathematics should be held as an "authoritative discipline" but the case for this is unclear and might rely on a move that Maddy rejects from the utility of mathematics to mathematical realism.

Returning to the broader question of the paper, if Thin Realism is false and Maddy shouldn't adopt descriptive criteria for axiom selection, a positive proposal of how to justify the sort of account that both Maddy and myself want is owed. This paper's positive proposal aims to provide just that.

# 3 An Instrumentalist Theory of Axiom Selection

I've made several promises about what my positive proposal will do and how it will improve on Maddy's position. For reference, here's a brief summary of what my account of axiom selection aims to do:

1. Provide an account that's extensionally similar to Maddy's.
2. Avoid the problem of aggregating individual goals into collective goals.
3. Explain why depth, or similar mathematical virtues, are valuable for inquiry.
4. Avoid the Euthyphro dilemma.
5. Avoid relying on Thin Realism.

I'll sketch the positive proposal in brief, then go into details. We abandon Maddy's Thin Realism in favour of Instrumentalism about mathematics (explained in §3.1). The Instrumentalist takes a lot of inspiration from Maddy's Arealist. In particular when Maddy says *"the Arealist holds that mathematics is distinguished from other extra-scientific enterprises by its role as a handmaiden to science"* (**SP** - IV.4, p386), she describes a sentiment that the Instrumentalist takes to heart. On the Instrumentalist view, the (main) purpose of Mathematics is as the handmaid of the sciences. The ultimate function of mathematics is to build formal tools for scientific inquiry. Algebra, geometry, calculus, group theory and more are all examples of tools mathematics has built over the years. These tools have been invaluably useful. The details of how these tools are used is discussed in §3.1.

The axiomatic method builds the sharpest (most precise and logically rigorous) tools. Axiom systems are evaluated on the basis of their efficacy in producing tools. Sometimes, they are introduced to define a single tool (e.g. the axioms of group theory). Sometimes, they provide a playground for mathematicians to tinker, build many tools and understand their consequences (e.g. Set Theory). New axioms are to be adopted if they facilitate the building of quality tools. This consideration generates a landscape of extrinsic reasons for adopting axioms. The kinds of structures that an axiomatic system entails are the metric by which it is evaluated.

An important consequence of the view, as advertised, is axiomatic pluralism. There is no reason for the Instrumentalist to limit themselves to a single axiomatic system. They are not laissez-faire about this, there are some axiom systems that are not productive nor worth the time of mathematicians. But they see no reason to decide indeterminate statements simply for the sake of deciding them. Interesting structures might exist in both branches extensions could take. This is discussed in §3.2.

Going through the aims for the account, Thin Realism is avoided by adopting an Instrumentalist account of mathematics. There is no need to demonstrate that my methods reliably produce true claims if truth isn't relevant to mathematics. The aggregation problem and the Euthyphro dilemma are similarly avoided because, unlike Maddy, I avoid an agent-based account of value in mathematics. If I can explain why depth, or similar virtues, are valuable for inquiry, then recovering an account extensionally similar to Maddy's should follow, given the discussion of §1.2. Mathematical virtues like depth are, I argue, an indicator that what one has is a well-developed tool. That doesn't necessarily mean that this tool has a real-world application but, if it does, it'll be an effective one. Interesting mathematical structures won't necessarily be applicable in a real-world case but, if it is applicable, it will reveal a lot about the phenomenon it's being used to model.

## 3.1 Mathematical Instrumentalism

### 3.1.1 The basic ideas...

Instrumentalism claims that the main purpose of mathematics is to build formal tools for the sciences. That this is *a* function of mathematics is not controversial but that it's the main function is explicitly rejected by, amongst others, Maddy.

This is also not to say that this is the only function mathematics plays. There may be other secondary functions. For the present argument to work, though, I need to at least claim that (1) tool building is the primary function of mathematics and (2) tool building is the function of mathematics relevant to axiom selection.

Moreover, this claim needs to be understood at the social, not the individual level. Mathematics, aside from being a body of thought, is a social institution. It is a collection of individuals, organizations, journals, archives, pedagogical, financial and administrative structures, and more. The Instrumentalist claims that the reason why we should have this kind of social structure is (primarily) because it builds formal tools that are exceptionally useful for science.

What are alternatives to Instrumentalism?

In the philosophy of science, Scientific Realism is often framed as the alternative to Scientific Instrumentalism. The same won't do here. As I discuss below, one can be a Mathematical Realist and nevertheless be an Instrumentalist.

Instrumentalism is a claim about why mathematics ought be pursued. It doesn't say anything about the nature of mathematical truth or objectivity. One thing that, incidentally, there are objective mathematical facts, but that the reason why it's worth doing mathematics has nothing to do with these facts.

The counterpoint to Instrumentalism, then, is not Realism but the view that mathematical facts, understood in a suitably Realist way, are worth knowing for their own sake and the primary function of mathematics is simply to uncover those facts, with secondary concern at best for their application. We might call this Descriptivism about mathematics.

### 3.1.2 Why be an Instrumentalist?

Sadly, a full justification of Instrumentalism is too large a task for this paper. All that's claimed here is that adopting Instrumentalism leads to an account of axiom selection which satisfies the success conditions outlined above. That being said, a few brief comments in favour of Instrumentalism.

If one is a nominalist, like myself, then one is most likely going to be an Instrumentalist. For nominalists, then there's not really much else one might take the purpose of mathematics to be. It couldn't be to uncover the mathematical facts, if no such facts exist.

But Instrumentalist is not simply a position for nominalists. As discussed in Essay II: *Mathematics Needs no (Philosophical) Foundation*, one might wish to hold that mathematicians should not need to make substantive metaphysical claims in order to adopt certain axioms. In that essay, I call this practical agnosticism. If one thinks this, one cannot be a Descriptivist about mathematics. For a Descriptivist, adopting an axiomatic theory would involve taking a stance on the nature of objective mathematical facts. But that would be a metaphysically substantial position, which mathematicians should be able to avoid if they wish.

Instrumentalism is a more appropriate view of Mathematics if one wishes to keep Mathematics neutral with respect to substantive philosophical questions.

All this being said, whilst arguments from Nominalist or Agnostic to Instrumentalism are clear, one can be nevertheless a realist and an Instrumentalist.

Suppose one thinks that the purpose of mathematics is to build tools for science. It remains to be seen if those tools work, how well they work and if they are essential or simply useful. The Instrumentalist-Realist is then surprised to find that they work exceptionally well, unreasonably well even. Moreover, science doesn't look like it could be done without these tools. On these grounds, and following Quine (1981)'s familiar reasoning[9], the Instrumentalist-Realist concludes that there must really be mathematical facts out there that play an essential role in the workings of the universe.

But the Instrumentalist-Realist is still an Instrumentalist. It's a potentially interesting conclusion that there really are mathematical objects, but learning about these objects was never the reason for doing Mathematics. That Quine has convinced them that mathematical objects really exist changes nothing about why they care about mathematics.

[9] See also **SP** IV.2 for a broad discussion not just of Indispensability Arguments but a broader array of arguments from the utility of Maths in science to Mathematical Realism.

The Instrumentalist-Realist, then, disagrees with Quine on axiom selection. Quine is a Descriptivist. For Quine, axiomatic theories are descriptive theories of the mathematical facts. They are motivated as a way of describing or explaining said facts. Quine is limited to adopting a relatively minimal system that merely explains the mathematics needed for science. The Instrumentalist-Realist disagrees thoroughly. The purpose of axiomatic systems is to provide a workspace for mathematicians to build and investigate kinds of tools. Consequently, they need not restrict their axioms to a comparatively minimal system for unifying and explaining the truth of certain parts of mathematics. They are realists about many parts of mathematics, but their realism is subsequent to and not a restriction on mathematical practice.

#### 3.1.3 What kinds of tools does mathematics produce? How is this useful for science?

The Instrumentalist claims that the function of mathematics is the production of formal tools for the sciences. But tools are an exceptionally diverse category. What kinds of tools does mathematics produce?

Mathematics produces tools for abstracted and idealized representational of physical phenomena[10]. Physical structures are represented by mathematical analogues that are formally or logically similar with respect to certain important features. Spacetime can be represented as a 4D coordinate system, for instance, because spatial and temporal relations are (at least in Newtonian cases) transitive, asymmetric, irreflexive, have four degrees of freedom, etc. The logical structure of spacetime is similar to the logical structure of the mathematics used to represent it.

Of course, idealization is made. Spacetime might be discrete, but at a small enough scale that it can be represented using continuous geometric spaces. Part of applying some mathematics to science is understanding the idealizations being made and if, or the extent to which, that might impact the outcome.

Representational tools are an important part of modelling. Modelling is an important part of making scientific predictions.

[10]My account is influenced both by Leng (2021)'s discussion of the explanatory role of mathematics and also Maddy's discussion in **SP** IV.2.

Now the nature and role of modelling in science is an intensely discussed topic within the philosophy of science. In particular, if all models are representational or how models justify is a hot topic. I certainly wish to avoid making any overly strong claims about what models are or what they do. Suffice for present purposes to make the following claims: (1) There are at least some, in fact many, scientific models that are representational (2) Some of them are representational because they are mathematical models wherein some of the mathematics is used to represent, with idealization, their intended target (3) Some of these still make inferences about their targets by making inferences about the mathematical structures used to represent the target and applying the result back by analogy.

There might be exceptions to this story but it's sufficient for present purposes that something sufficiently like this happens frequently.

A nice, though well-worked, example of how mathematics can be used as a representational tool is the Lotka–Volterra model of population dynamics. The model has three variables (time, predator-population, prey population), two differential equations governing the relationship between these variables and some initial conditions.

The differential equations link change in prey population over time to predator size, and vice versa. Let *prey* be the variable for the prey population and *pred* for the predator.

$$\frac{d - prey}{dt} \propto pred^{-1} \qquad\qquad \frac{d - pred}{dt} \propto prey$$

It's very clear how the mathematics in this model is representational. The variables have clear real-world interpretations and the equations represent a causal relationship between the three quantities at play. The initial conditions will represent some contingent facts about a case being studied.

Idealizing assumptions are made here. Population is discrete, not continuous. Predator-prey interactions are probabilistic in nature, meaning there will be some random noise. Some prey can also kill their predators. There are external conditions such as the seasons (perhaps the prey is better at hiding in winter) that are not factored into the model. Both the predators and prey have evolutionary pressures that will change the coefficients of the differential equations.

Nevertheless, there are sufficient similarities between the mathematical model and the physical system it represents that some reasonable inferences can be drawn, within sensible bounds.

I have picked a case friendly to the Instrumentalist here. Lotka-Volterra model is a nice basic example, but it is simplistic. The way in which more complex models represent is far more abstract. Something like a neural net might end up with a highly non-representational internal structure (e.g. where the individual node values have no clear interpretation). One might imagine an AI weather model that takes real-world data as an input, processes this in a way that's not humanly understandable, and then outputs weather probabilities. This would only be representational at the input and output stage and not representational through the middle. Its representational nature should be understood holistically, at the level of the structure, not necessarily atomistically at the level of the components of the model.

There's certainly a lot more to say here. Scientific models are incredibly diverse and I'm certainly not saying that how mathematical representation works in the case discussed is how it will work in all cases. I don't wish to over-extrapolate here. But my hope is that upon inspecting a wider range of models a similar relationship between mathematics and physical structures, perhaps realized in a different way, will be found.

As a final tangential comment here, one can see here a relationship between Structuralism about Mathematics and Instrumentalism. I understand Structuralism as the claim that mathematical statements do not have particular content, but rather are generic or structural in nature. $\in$ is not some *particular* binary relation satisfying the axioms of set theory but is a stand-in for any relation that happens to satisfy that property. Similarly, $<$ is not some particular ordering relation but a stand-in for any order.

If Instrumentalism is true, it's clear why one might want Mathematics to be structural. Instrumentalists value mathematics because of its utility as a representational tool. They want to be able to apply these tools as widely as possible. By having structural, rather than particular, content, mathematics can be applied to any target phenomena which, modulo sufficient idealization, are particular instances of general structures.

In summary, modelling is an incredibly important part of knowledge generation within science. Mathematics is incredibly useful for modelling so mathematics should be pursued as a means of developing the kinds of representational tools needed for science. So claims the Instrumentalist, this is the reason why mathematics, as a whole, is worth doing.

#### 3.1.4 Returning to axiom selection

A naive Instrumentalist position would simply say that axiomatics should follow the needs of science. An axiomatic theory is good iff it defines or in some other way helps produce mathematical structures which are useful for the modelling or representational needs of the science of the day.

This would be, I think, more than a little short-sighted. Perhaps paradoxically, it seems as if the best way to get mathematicians to produce useful tools for science is to give them some degree of free reign to pursue the mathematics that they find interesting.

There's something also to be said for letting a thousand flowers bloom. Let the pure mathematicians build and explore scores of formal tools to their heart's content, then let the applied mathematicians or scientists pick from and adapt these where appropriate. A particular tool might look mathematically interesting and be worth investigating on these grounds but only prove its value in application later.

That's not to say that any mathematics goes. As Maddy discusses (**DTA** II) there are criteria that mathematicians apply when evaluating how interesting a piece of mathematics is. There are certain virtues that make particular structures more or less interesting (where "interesting" should here be read as something like "worthy of investigation").

What I want to suggest is a degree of removal on the part of the Instrumentalist. A series of mathematical virtues (depth, simplicity, novelty, etc) are identified. In the long run, the Instrumentalist holds that these virtues are good indicators that a bit of mathematics will serve as a useful representational tool. But (pure) mathematicians don't need to worry about possible future applications. They just need to continue as they are, guided by these virtues. In the long run, then, useful tools will be produced.

To slightly clarify the above statement, for at least some of the virtues, it's not that they necessarily indicate a higher likelihood that a mathematical structure will be useful in application but rather that if such an application is found, it'll be a very useful one which reveals interesting and unexpected features of the structure being modelled. Maddy calls this the "more out than in" miracle (**SP** IV.2).

How does this relate to axiom selection?

Clearly, the Instrumentalist wants axiomatic theories which maximize these virtues, but this is a slightly over-simplistic picture. Different sorts of axiomatic theories are used in different ways. It's worth drawing a distinction between two types of cases.

In some cases, axioms are used to define a singular kind of structure. Group theory, for instance, defines a particular sort of structure - the group. The same is true of something like arithmetic as well, the axioms of which define number-systems. These axiomatic theories are akin to singular tools (or perhaps tool types). They represent a frequently occurring or especially interesting type of pattern or structure. In this case, the Instrumentalist simply wants the structure described by the axioms to possess the kinds of virtues mentioned above.

To continue with the example of group theory, there are multiple ways in which groups instantiate the mathematical virtues. One of the nicest examples is the way in which it connects seemingly distant parts of mathematics. For instance, it's used to study symmetry and transformations on ordinary objects but, via Noether's theorem, is also deeply connected to conversation laws in mechanics. Finding unexpected connections between different parts of mathematics is an archetypal example of mathematical depth.

In other cases, axiomatic theories do not necessarily define a singular structure which is itself useful but rather is used to build and analyze a range of different structures. Set theory is the archetypal example of this. It's tempting to think of these types of axiomatic theories like toolboxes: their function is simply to contain as many useful tools as possible. But this is not the whole picture.

It's true that a useful feature of set theory is that it allows one to build a wide range of mathematical structures, many of which instantiate the virtues mentioned. But this is not its only function. Set theory can also be used to analyze the structures they contain. There are facts about $\mathbb{N}$ which can be proven in $ZF$ but not in $PA$.

So these kinds of "toolbox" theories (I'm deliberately avoiding the term "foundational") don't just contain a number of interesting structures which possess the kinds of virtues discussed, they also help better understand those structures.

These two categories are not mutually exclusive. Geometry, for instance, arguably does a little of both. Geometric spaces are, themselves, interesting structures but contain things like triangles or other shapes which are (obviously) useful structures. Arguably models of $ZF$ are also, in and of themselves, interesting structures and $ZF$ is a good axiomatic theory not simply because of its "toolbox" role but also in the more direct sense, like in the case of Group Theory.

Thus, granting the Instrumentalist's claim that the kinds of mathematical virtues mentioned do, in fact, indicate useful applicability, axiomatic theories are to be evaluated on three basies: (1) their ability to directly define structures that, themselves, possess the virtues (2) their ability to provide a "toolbox" of interesting structures possessing the virtues and (3) their ability to contribute towards investigations of the aforementioned structures contained within the "toolbox".

The only thing left to do, then, is to justify the Instrumentalist's claim about the mathematical virtues.

### 3.1.5 On Depth

I don't attempt to offer either a list of all the possible mathematical virtues or defend each of their connections to the possibility of useful application. I expect there will be a little heterogeneity here. Different virtues might connect to useful application in different ways, and even the way in which a particular representation might be useful might vary from case to case. I certainly don't wish to claim that the account of the value of depth given here can be applied to other cases with little change. Perhaps it can, but that's left to future work.

As proof of principle, this section outlines an Instrumentalist account of the value of mathematical depth. In Maddy's account of axiom selection, depth plays a central role as one of the primary goals of mathematical investigation. Given my above-stated aim of recovering an extensionally-similar account of axiom selection to Maddy's, demonstrating the value of depth is important.

Even still, I must further clarify exactly what I achieve here. Since **DTA** there has been a moderate but thorough discussion on the nature of mathematical depth. Arana (2015), Lange (2015), Stillwell (2015) and Urquhart (2015) together formed a special issue of *Philosophia Mathematica* on the topic. See also Weisgerber (2023) for a more recent discussion in particular of Maddy's account. I only engage in the concept of depth as it appears in **SP** and **DTA**.

This restriction is justified on two grounds. (1) As I'm trying to recover an account of axiom selection extensionally similar to Maddy's, it's Maddy's notion of depth that matters (2) this section is in part intended more as proof of principle than a complete and final account of the Instrumentalist's theory of axiom selection. Reappraising depth (and other mathematical virtues) would form part of future work to turn this promissory note into a fully worked-out account.

Clarifications aside, what value does mathematical depth have, according to the Instrumentalist?

Depth, in Maddy's sense, is not per se a property of mathematical structures but rather involves the discovery of unexpected or novel connections between typically distant parts of mathematics.

The example discussed above of the relationship between conservation laws and geometry symmetry via groups is, I think, a clear example, though not the one that Maddy gives.

An example that appears in both **SP** and **DTA**, also discussed above, is the relationship between large cardinals and the claim that the axiom of determinacy holds in the smallest inner model of $ZF$ containing the reals ($AD^{L(\mathbb{R})}$). Large cardinals imply $AD^{L(\mathbb{R})}$.

It's surprising that the structure of $L(\mathbb{R})$ would be related to the presence of large cardinals. That such a connection exists is an example of depth.

Why would unexpected connections indicate useful application?

As noted above, I'm not claiming that the presence of deep connections makes it more likely that some mathematics has any application. What I'm claiming is that, conditional on there being any application at all, this application is more likely to be useful.

Because depth involves unexpected connections between seemingly separate parts of mathematics the discovery of depth reveals a lot of information about the connected structures. It allows for the application of different methods and theorems not previously known to be applicable. The solution to Fermat's Last Theorem, for instance, was only possible because it could be re-described in terms of elliptical curves. The discovery of a deep connection goes hand-in-hand with new discoveries and understanding of the target structure.

A mathematical structure deeply connected to other areas of mathematics is likely, therefore, to (1) be comparatively well understood and (2) have unexpected or non-obvious features.

If a physical system can be well represented by just such a structure, the analogous facts hold. Via the modelling process, and assuming that any idealizing assumptions do not get in the way, a potentially large number of unexpected or non-obvious conclusions can be drawn.

In **SP** IV.2, for instance, Maddy discusses what she called the "more out than in" phenomenon, the phenomenon where the application of well-chosen mathematics leads to new physical discoveries. A nice example she gives is that of the Dirac equations in Quantum Mechanics.

The Dirac equations came out of the desire to reconcile relativistic time with a quantum mechanical account of electromagnetism. This was not a purely mathematical concern but it turns out that the mathematics was possible only if at each spacetime point there are four wave functions, not two. This entails the existence of particles with the same mass as electrons but opposite charge (positrons). These particles were theorised as a consequence of the mathematics used but only discovered later.

Here surprising and novel features of the mathematics (the required number of wave functions at each point) lead to a discovery about the nature of the physical systems being modelled.

Depth then seems to display the desired features. A mathematical structure with deep connections to other parts of mathematics, when successfully applied, yields surprising new information about the target structure. Depth, then, is a mathematical virtue because it helps mathematics fulfil its function of producing formal representational tools for modelling that are useful for science.

### 3.1.6 Taking stock...

Five things were promised of the Instrumentalist's account:

1. An account that's extensionally similar to Maddy's.
2. That avoids the problem of aggregating individual goals into collective goals.
3. That explains why depth, or similar mathematical virtues, are valuable for inquiry.
4. That avoids the Euthyphro dilemma.
5. That avoids relying on Thin Realism.

(1) and (3) follow from the previous two sections. The value of the kinds of mathematical virtues that are the objects of the (more fundamental) goals of mathematicians can be explained on the Instrumentalist's account. These virtues are valuable because mathematical structures instantiating them are, if successfully applied, more informative and hence useful tools. In the case of depth, they reveal a great deal about the target physical structure.

(2) and (4) are met by avoiding any reliance on the goals of mathematicians. Instead, I rely on the function of mathematics as the guiding normative principle[11].

(5) is met due to the theory-neutrality of Instrumentalism, as discussed in §3.1.2. One can be a Realist, Agnostic or Nominalist whilst still being an Instrumentalist.

[11] In meta-ethical terms, I'm using a telic rather than an agential theory of (mathematical) value,

I've shown, I hope, that an Instrumentalist theory of axiom selection does all that was advertised of it. It retains the positive parts of Maddy's theory (its functionality in application) whilst fixing many of its theoretical issues.

## 3.2 Corollary: Axiomatic Pluralism

There is, I think, one difference between my Instrumentalist account of axiom selection and Maddy's account: Axiomatic Pluralism[12].

The Instrumentalist account of axiom selection entails that there are many correct axiomatic theories, not simply one (relative to particular domains).

Maddy claims that one of the goals of set theory is to describe a singular set-theoretic universe[13]. Maddy might be wrong about this. One could definitely read this section not as outlining a difference between my Instrumentalist and Maddy but as explaining, on Maddy's terms, why she should be an Axiomatic Pluralist.

As is, though, I'll take it that Maddy has a better understanding of the consequences of her own position than I do and present this as a point of difference between the two views.

### 3.2.1 Instrumentalism entails Axiomatic Pluralism

To see why the Instrumentalist is an Axiomatic Pluralist, consider cases where both some prospective axiom $\phi$ and its negation have interesting consequences. That is, to put it more exactly, they both entail the existence of structures either of practical application or containing the kinds of hallmark mathematical virtues that indicate it might one day be of use (depth, interest, etc). The two might not be interesting in equal measure, but both are interesting.

[12]Axiomatic Pluralism is not a position that has been extensively discussed or defended in the literature. Michèle Friend has discussed "foundational pluralism" as part of her extensive discussion of pluralism in (the philosophy of) Mathematics (Friend (2013) and Friend (2019)). Axiomatic Pluralism on the basis of Logical Pluralism has been discussed by Davies (2005), Hellman and Bell (2006), Koellner (2009) and Priest (2019).

[13]Maddy's view might have changed on this given recent work by her and Toby Meadows (Maddy and Meadows, 2020). At the least, her view at the time of **SP** and **DTA** was a monist.

Something like $CH$ is a good example of this. $\neg CH$ models are often more interesting. There are frequently properties of cardinals bound that are bigger than $\aleph_0$ but are at most $2^{\aleph_0}$. There might be some property $P$ such that $\neg P(\aleph_0)$ but $P(2^{\aleph_0})$. The interesting mathematical question is what the smallest $P$ cardinal is. Or if the smallest $P$ cardinal must be greater or smaller than the smallest cardinal satisfying some other similarly bound property.

In $CH$ models these all just collapse upwards into $2^{\aleph_0}$ but $\neg CH$ models allow for far more versatility. But this doesn't mean that $CH$ models aren't also interesting.

Because Maddy wants a decisive set theory, she is forced to choose. She must throw out interesting structures whichever option she takes. Naturally, she takes the lesser of two evils and keeps the more interesting structures in, but she is forced to pick some evil.

The Instrumentalist, on the other hand, sees no need to choose. They consider the consequences of both $CH$ and $\neg CH$ and flit between them for different purposes. Interesting tools can be built in both kinds of model. For this reason, both $ZFC + CH$ and $ZFC + \neg CH$ are correct systems to adopt.

In short, Axiomatic Monists, like Maddy in **DTA**, will adopt an axiom iff it is better qua their evaluation metrics than its negation. This forces them to throw out interesting structures. Axiomatic Pluralists, like the Instrumentalist, will only accept an axiom if its negation fails to yield any interesting structures. In the case where both an axiom and its negation are interesting, the Monist chooses and the Pluralist takes up both.

# Conclusion

This paper outlines an alternative to Maddy's account of axiom selection based on Instrumentalism about mathematics. §1 describes Maddy's view in *Second Philosophy* (Maddy, 2007) and *Defending the Axioms* (Maddy, 2011). Two issues are identified during §1 but the more substantial issues are identified during §2: the Euthyphro Dilemma and my objections to Thin Realism. §3 then presents my positive account, showing how it avoids the issues of Maddy's position outlined in §1 and §2, but keeps the virtues of the account. §3.2 then outlines a (potential) difference between my position and Maddy's, namely Axiomatic Pluralism.

## References

Arana, A. (2015). On the depth of szemeredi's theorem. *Philosophia Mathematica*, 23(2):163–176.

Davies, E. B. (2005). A defence of mathematical pluralism. *Philosophia Mathematica*, 13(3):252–276.

Feferman, S. (1999). Does mathematics need new axioms? *American Mathematical Monthly*, 106:99–111.

Feferman, S. (2000). Why the programs for new axioms need to be questioned. In Solomon Feferman, Harvey Friedman, P. M. and Steel, J., editors, *Does Mathematics Need New Axioms?*, volume 6, pages 401–446. Bulletin of Symbolic Logic.

Friend, M. (2013). *Pluralism in Mathematics: A New Position in Philosophy of Mathematics*. Dordrecht, Netherland: Springer.

Friend, M. (2019). Varieties of pluralism and objectivity in mathematics. In Centrone, S., Kant, D., and Sarikaya, D., editors, *Reflections on the Foundations of Mathematics: Univalent Foundations, Set Theory and General Thoughts*, pages 345–362. Springer Verlag.

Hellman, G. and Bell, J. L. (2006). Pluralism and the foundations of mathematics.

Koellner, P. (2009). Truth in mathematics: The question of pluralism. In Bueno, O. and Linnebo, , editors, *New Waves in Philosophy of Mathematics*, New Waves in Philosophy. Palgrave Macmillan, London.

Lange, M. (2015). Depth and explanation in mathematics. *Philosophia Mathematica*, 23(2):196–214.

Leng, M. (2021). Models, structures, and the explanatory role of mathematics in empirical science. *Synthese*, 199(3-4):10415–10440.

Maddy, P. (1993). Does v. equal l? *Journal of Symbolic Logic*, 58(1):15–41.

Maddy, P. (2007). *Second Philosophy: A Naturalistic Method*. New York: Oxford University Press.

Maddy, P. (2011). *Defending the Axioms: On the Philosophical Foundations of Set Theory*. Oxford, England: Oxford University Press.

Maddy, P. (2019). What do we want a foundation to do? In Centrone, S., Kant, D., and Sarikaya, D., editors, *Reflections on the Foundations of Mathematics: Univalent Foundations, Set Theory and General Thoughts*, pages 293–311. Springer Verlag.

Maddy, P. and Meadows, T. (2020). A reconstruction of steel?s multiverse project. *Bulletin of Symbolic Logic*, 26(2):118–169.

Moore, G. H. (1982). *Zermelo's axiom of choice : its origins, development, and influence*. Studies in the history of mathematics and physical sciences ; 8. Springer, New York, NY [u.a.].

Moriarty, C. M. (2022). Ructions over fluxions: Maclaurin's draft, the analyst controversy and berkeley's anti-mathematical philosophy. *Studies in History and Philosophy of Science*, 96:77–86.

Moriarty, C. M. (2023). Duelling catechisms: Berkeley trolls walton on fluxions and faith. *Intellectual History Review*, 33(2):205–226.

Priest, G. (2019). From the foundations of mathematics to mathematical pluralism. In Centrone, S., Kant, D., and Sarikaya, D., editors, *Reflections on the Foundations of Mathematics: Univalent Foundations, Set Theory and General Thoughts*, pages 363–380. Springer Verlag.

Quine, W. V. (1981). *Theories and Things*. Cambridge, Mass.: Harvard University Press.

Rosen, G. (1999). Review of maddy [1997]. *The British Journal for the Philosophy of Science*.

Schlimm, D. (2013). Axioms in mathematical practice. *Philosophia Mathematica*, 21(1):37–92.

Stillwell, J. (2015). What does 'depth' mean in mathematics? *Philosophia Mathematica*, 23(2):215–232.

Urquhart, A. (2015). Mathematical depth. *Philosophia Mathematica*, 23(2):233–241.

Weisgerber, S. (2023). Mathematical progress — on maddy and beyond. *Philosophia Mathematica*, 31(1):1–28.

# Essay IV: Proofs, Derivations and Axiomatic Pluralism

Recent work in the philosophy of mathematical practice by Azzouni (2004), Burgess and Toffoli (2022) Toffoli (2021), Hamami (Hamami (2014) and Hamami (2022)), Hamami and Morris (ming) and Tanswell (Tanswell (2015) and Tanswell (2016)) and others[1] has explored the relationship between mathematical proofs and derivations.

The term "proof" should be here understood as the kinds of things one finds in mathematics journals or books; written in natural languages, though perhaps with some formalization where appropriate, perhaps also containing diagrams (Toffoli, ming) or other kinds of non-formal reasoning. The term "derivations" should be understood as a logically valid sequence of statements, written in a formal language, proceeding from some choice of axiomatic theory; the kinds of things many of us ask our students to produce in our introductory logic classes.

Trends in this work support a greater gap between proofs and derivations than what was previously been thought. Proofs are *indicators* of the existence of a derivation (Azzouni, 2004), but not strict instructions for producing one.

That a proof can be faithfully unpacked into multiple different derivations is not a new claim within this literature. What is new, and what I draw attention to here, is that proofs can be faithfully unpacked into multiple different derivations *proceeding from different axiomatic theories*.

[1] See also Kitcher (1981), Fallis (2003), Detlefsen (2008), Leitgeb (2009), Marfori (2010) and Andersen (2020)

If this is true, then there are potential ramifications for the philosophy of axiom selection. I argue that it makes a kind of axiomatic pluralism more plausible.

§1 clarifies some definitions. §2 discusses the relationship between proofs and derivations. §3 considers ramifications for the philosophy of axiom selection.

# 1 Some definitions

As mentioned above I use the term "proof" to indicate the kinds of things that mathematicians write that appear in mathematics journals. *Modular elliptic curves and Fermat's Last Theorem* (Wiles, 1995) or *The Independence of the Continuum Hypothesis* (Cohen, 1963) are both examples of proofs. For more examples, one may open the latest edition of *Annals of Mathematics* and encounter dozens. Finding the best definition proof in this sense is an open problem in the philosophy of mathematical practice. Azzouni (2004) is right that they are derivation indicators, but this is a necessary not a sufficient condition. Suppose I read an article in a reliable news outlet reporting on a recently published proof. That is also an indicator of a derivation but is not a proof. Suffice for present purposes to say that proofs typically have the following features:

1. They are typically written in natural languages, often English, likely augmented with technical vocabulary and some formalism where appropriate.

2. They frequently have the form of a series of declarative statements (lemmas and theorems) connected by instructions in the imperative mood for navigating between them. E.g. "Let A and B be spheres. Construct a line C tangential to both A and B..."

3. Steps of the proof are not immediate logical consequences of one another. It requires a level of mathematical understanding to connect the steps of a proof.

4. Proofs infrequently explicitly state their axioms.

5. Proofs might contain diagrammatic elements.

This is to be contrasted with derivations, which are called "proofs" in most introductory logic books that philosophically educated readers are more likely to be familiar with.

Derivations are a sequence of declarative statements, written in a formal language and constrained in their order by certain logical inference rules. There might be certain statements allowed in a proof as a result of them being logical axioms. Other statements that appear in a derivation which neither follow from previous statements via inference rules nor are logical axioms are the mathematical axioms (or just axioms) of the derivation.

I will use the verb "unpack" to relate a proof and a derivation that are related in the appropriate way. Hales' *"A proof of the Kepler conjecture"* (Hales, 2005) is a proof which unpacks to the derivation presented in *"A formal proof of the Kepler conjecture"* (Hales et al., 2017) as produced by the Flyspeck project.

Whilst not every derivation from some axioms to the conclusion of a proof is an unpacking of that proof, I'll be somewhat liberal in my use of the term "unpack". Any derivation that could at all reasonably construed as a rational reconstruction of the derivation indicated by a proof can be considered an unpacking of that proof.

Of course, there are a number of normative standards one might apply when relating a proof to a derivation. A derivation might be more or less faithful an unpacking of a proof based on how well it captures some combination of (1) what is written in the proof (2) what the author(s) had in mind and (3) the social-epistemic context in which the proof is produced.

(3) is important to note. Proofs do not exist in a social vacuum. They rely on theorems, definitions, ideas and methods present in the social-epistemic context in which are presented.

This paper's main claim should therefore be clarified. It is not simply enough to show that there are proofs which unpack to derivations with different axioms, that is easy to show. Rather, I wish to claim that there are proofs which *faithfully* unpack to multiple derivations with different axioms. The question is then what the appropriate notion of faithfulness is.

## 2 How to unpack proofs into derivations

### 2.1 The relationship between proofs and derivations

I want to think of the relationship between proofs and derivations a little like the relationship between higher-order programming languages and machine code[2]. In a computer, a CPU takes instructions in some form of machine code. This is a highly basic programming language which would be utterly impractical to code in on a day-to-day basis. In practice, programmers code in higher-level programming languages such as Python, Ruby, Java and C#. These are languages which allow for more abstracted, simple descriptions of processes which might be quite complex at the level of machine code. This allows programmers to simply get on with coding without needing to worry about orthogonal base-level issues.

Higher-level code is connected to machine code via a compiler[3]. Compilers are effectively an algorithm for taking instructions written in a higher-level language and translating them into machine code, which the processor can read. The majority (but not all) of this amounts to a series of definitions connecting procedures in the higher-level language to procedures in the lower-level language.

For instance, take the following simple Python script:

```
if a*b==c:
    print(str(c))
```

This script compares three variables - a,b and c. If a times b is c, then the computer prints the value of c.

This is easy enough to understand even if one has never encountered the language of Python before. A compiler would take this and turn it into something like the following. I describe it in English as neither I nor the vast majority of my readers can read ML code.

[2]Yacin Hamami should be credited for this particular analogy. I had the fortune of attending his work-in-progress talk whilst working on an earlier version of this paper and his analogy with programming languages has proven a useful way of framing their relationship.

[3]For this paper, I ignore the distinction between compilers, interpreters, IDEs, etc. For present purposes, the term "compiler" is used in its more generic social meaning.

Take the memory addresses of a and b. Going to those, perform a multiplication algorithm, saving the result at a third memory address. Compare that memory address with c. If they are the same, go to a fourth memory address and write the code of the text representation of the characters in the value of c. Send instructions to the monitor to display the text stored at the fourth memory address.

A compiler is then, effectively, a translator between the higher-level languages of day-to-day coding and machine code.

A similar story can be given of the relation between proofs and derivations.

Proofs are written in augmented versions of natural languages. New specialised words are added to the base language and some ordinary words are redefined in more precise manners (e.g. consider the difference between the ordinary English and Mathematical senses of "similar"). Derivations are written in formal languages. Just as a processor couldn't understand instructions fed to it in Python or C#, someone only familiar with the formal languages used in derivations would not be able to understand a mathematical proof.

In order to connect proofs to derivations, something must play the role of the compiler.

One part of this is (relatively) easy. There are usual and known methods for translating natural languages into formal languages. We have our undergraduate students do this as part of their logic exams all the time! Perhaps the only challenging element consists in translating the imperatives frequently found in mathematical proofs into the declaratives of formal languages. But this is relatively unproblematic when one sees that imperatives are typically higher-level descriptions of which inference rules to apply to the statements one already has, rather than declarative statements in disguise.

The more challenging part of unpacking a proof into a derivation comes in translating the higher-level concepts of a proof into the lower-level concepts of the derivation. For example, consider the concept of prime factorization. This is a common enough mathematical concept, but one that cannot be (directly) found in the formal languages used in derivations. The language of arithmetic, for instance, contains things like the successor relation, the functions $+$ and $\times$, etc. The higher-level concept must be defined from lower-level building blocks.

For instance, a prime factorization of X can be defined as a factorization of X where all the elements are prime. A factorization of X can be defined as a sequence of natural numbers such that their product is X. A prime number is a number whose only factorizations are sequences containing itself and arbitrarily many 1s. Sequences might then be defined in set theory as ordered tuples, which in turn are defined from sets. Natural numbers can be defined in set theory in familiar ways.

Three things to note here.

First, there's a degree of layering present here (and this is again similar to programming languages). There are not merely two levels of language at play: the formal language of derivations and the non-formal language of proofs. Going "up" from the language of derivations, there are increasingly abstract concepts defined from previous concepts. In the above definition of a prime factorization, there are six levels of definition: sets, tuples and numbers, sequences, factorizations, prime numbers, and prime factorizations.

Second, there's no unique way of defining higher-level concepts. For one, the notion of a sequence is redundant in the above definition of prime factorization. I could have used the term "tuple" from the start and saved an entire definitional stage. But more radically, I could have used alternative, perhaps non-set theoretic, definitions of the higher-level concepts.

I could have defined a sequence of natural numbers as a function from an initial segment of $\mathbb{N}$ into $\mathbb{N}$. I could have defined a factorization of X not as a sequence but as a $f$ function from $\mathbb{N}$ into $\mathbb{N} \cup \{0\}$ such that $x = \prod_{x \in \mathbb{N}} x^{f(x)}$. Similarly, I could have avoided set theory and used higher order arithmetic to define the otherwise set-theoretic notions, though leaving $\mathbb{N}$ as basic. To put this in terms of the analogy, there are multiple possible compilers which unpack proofs into different derivations.

There is a point of intersection here with mathematical structuralism. When providing a lower-level definition of a higher-level mathematical concept all that matters is the relevant inferential features of the higher-level concept is preserved. One does not need to find the *real* meaning of, say, a sequence (whatever that would even mean). It suffices to find something definable in one's base language which *behaves* like a sequence behaves.

None of what's been said here breaks from the programming analogy. Programming languages have a more complex hierarchical structure than the one presented here and, even within a program written in a particular language, there will often be additional processes defined within the code. Moreover, there are frequently new, different or better compilers produced which can unpack higher-level code either (1) more efficiently or (2) into more efficient ML.

The last point of note interacts with the question of pluralism. In the case above, the different ways that the concept of "prime factorization" unpacks into a more fundamental concept changes the axioms that would be required, were the concept used in a proof. Three different definitions were suggested: the set-theoretic definition given in the example, the functional definition and the higher-order arithmetic definition. If one wished to derive something about prime factorization, e.g. say its uniqueness, the axioms required would be very different based on the way in which one represents the higher-level concepts present in the proof. In the first case set theoretic axioms would be required. In the second perhaps some type-theoretic axioms. In the third something like finite order arithmetic.

In a sense, the choice of axioms doesn't really have much to do with the proof of the theorem itself. All the important features of the proof take place at the level of higher-order concepts. All the axioms do, in this case, is provide a mechanism for making the higher-order concepts rigorous. It matters that there is some possible choice of axioms which lead to a derivation of the theorem; but, in this case, it doesn't matter all that much which.

This is then a version of the type of pluralism advertised. Any ordinary proof unpacks into multiple derivations. Based on one's representational choices, this includes derivations from different axioms.

This is perhaps a little unsurprising. I could have shown this simply by pointing to examples. Take Thomas Hales' proof of the Kepler Conjecture. Due to some controversy surrounding the status of Hales' then-purported proof, Hales set out to use new innovations in computing to produce a formal derivation of his proof. This was the Flyspeck project. The project was successful and a formal derivation of the Kepler conjecture is now available.

Flyspeck used a combination of the digital proof finders HoL Light and Isabelle, with HoL Light forming the basis and background of the derivation and Isabelle performing some heavy combinatorial work over a class of possible counterexamples to the Conjecture, eventually imported back into HoL Light. HoL Light is of interest here as it covered the foundational portion of the proof. The axiomatic theory used was a version of Church's Type Theory, augmented with axioms of choice and functional extensionality (Call this *CTT+*). The derivation was not a derivation from the axioms of standard set theory: Zermelo Frankel with Choice (*ZFC*).

Now the reasons for this were, to my understanding, purely computational. Type theories work well from a computational perspective. *CTT+* is a very practical, functional choice of axiomatic theory, in this context. Hales has not, to my knowledge, expressed any deeper philosophical support for *CTT+*. This appears to be a very utilitarian choice.

That being said, there's no principled reason why a derivation couldn't have been found from the axioms of *ZFC*. In fact, it's relatively easy to interpret Church's Type theory in *ZF*. Functional extensionality and choice follow from *ZFC*. It would only be a matter of adding a few more definitions and the derivation from *CTT+* could be converted into one from *ZFC*.

Thus there are clear examples of major proofs which can be unpacked into derivations with different axioms. What this section does, I hope, is explains *why* this happens. When providing lower-level formulations of higher-level concepts, one only needs to preserve the appropriate formal or inferential structure. But there are multiple ways of doing this in different kinds of formal languages with different axioms. This is comparable to how there are many compilers of higher-level computer code which produce different ML code given the same input. Thus ordinary proofs unpack to multiple derivations with different axioms.

That being said, as discussed in §1, different derivations might be more or less faithful unpackings of a particular proof. The interesting pluralist claim isn't that ordinary proofs unpack into multiple derivations with different axioms, rather it's the claim that there are multiple *faithfull* unpackings into derivations from different axiomatic theories.

Put differently, one might think of derivations as rational reconstructions of what goes on in a proof. A faithful unpacking is a good rational reconstruction. The interesting pluralist claim is that proofs can be rationally reconstructed using different axiomatic starting points.

Of course, the open question here is what the relevant understanding of "faithful" is. Is there a way of identifying one particular derivation as the privileged derivation?

In §2.2 and §2.3 I consider a range of ways of understanding "faithfulness", showing that each is either unattractive or leads to the kind of derivation-pluralism discussed.

## 2.2 The Bibliographical approach

Continuing with the example above, whilst both the unpacking of Hales' proof into a derivation from *CCT+* and into a derivation from *ZFC* might be perfectly reasonable unpacking, one of these has the advantage of being the derivation that, in a certain sense, was already "out there". One could take Hales' proof and do something like the following: (1) replace any (essential) citations with the relevant parts of the text they are citing (2) repeat until there are no more citations (i.e. take the transitive closure of the "cites" relation) (3) Go through and convert the resulting extensive natural-language text into a relevant formal language which formalises the claims and axioms of the proof. The resulting string of claims is, hopefully, a valid derivation. The axioms on which that derivation is based are then the *actual* axioms of the proof. Call this the bibliographical approach.

It's worth noting that this is very much not what Flyspeck did. Flyspeck involved a great deal of clarifying and re-working definitions and foundations to fit the demands of HoL Light and *CTT+*. It's likely, though I have not checked the transitive closure of Hales' original proof, the resulting derivation would be Set Theoretic, carving a path through known definitional work by Bourbaki.

There's a prima facie compelling reason to think of derivation picked out by the bibliographical approach as the privileged. After all, it is quite literally tracking the work on which a proof was based.

There are, I think, three reasons to be sceptical of this approach.

First, I think it's quite likely that, in practice, this approach would result in very hybridised derivations, using broad and inefficient axiomatisations. For example, suppose a proof both cites relevant work by Peano and Bourbaki. The resulting derivation, taken from the transitive closure of the citations, would contain both the axioms of Peano Arithmetic and *ZFC*. This would then be an unnecessarily baroque derivation, given that it's possible to define $\mathbb{N}$ in a model of *ZFC* and prove the validity of the axioms of *PA* for that structure.

A more faithful unpacking would involve ignoring what the citation history actually says, in favour of "cleaning up". Given that Boubaki has given us a way of avoiding commitment to the axioms of *PA* in addition to the axioms of *ZFC*, this should be taken advantage of. Finding the correct unpacking of a proof involves some level of idealization.

The second issue is that many published proofs contain mistakes. Of course, there are examples of *major* mistakes, faux-theorems that were accepted for a non-trivial period of time before being shown to be false (e.g. Goldfarb (1984) refuted a result published by Gödel (1933) some 50 years prior). But what I have in mind are more minor mistakes. Small inferences that don't seriously undermine the proof, the kinds of things that a mathematician knowledgeable in a particular area likely knows about, but also knows how to work around.

In practice, this is no great worry. Mathematicians know how to work around these kinds of errors, whilst taking away the relevant, more abstract, methodological insights. The idea or broad structure behind a proof is often more important than the literal words that were written.

But when applying the bibliographical approach, these mistakes are not filtered out. Again, the "correct" derivation is found by applying some degree of idealization and know-how.

My final objection to the bibliographical approach is from a particular case study: Andrew Wiles (1995)'s proof of Fermat's Last Theorem, his use of Grothendieck universes and McLarty (2010)'s "new" proof of Wiles' result.

A quick summary of the case: in the summer of 1993 Andrew Wiles proved Fermat's Last Theorem, a result eluding over three centuries of mathematicians. The proof was accepted and Andrew Wiles won countless awards and even Royal patronage for his work. Interestingly, from a foundational perspective, Wiles makes use of a type of mathematical structure called Grothendieck universes (henceforth, simply "universes). The existence of universes is not provable in *ZFC* as it implies the existence of a strongly inaccessible cardinal.

If one applied the bibliographical approach to Wiles' proof, one would, even modulo the aforementioned issues, reach a derivation from the stronger Tarski-Grothendieck Set Theory, not *ZFC*.

This was discovered by Colin McLarty who also showed how to complete Wiles' proof without using the full apparatus of universes. Now what McLarty did was, I think, rather interesting. He didn't take Wiles' proof and provide a workaround that avoided the use of universes. Instead, he showed how a weaker *ZFC*–friendly definition of universes could perform all the work that Wiles needed universes to do.

McLarty's amendment to Wiles' proof really amounts to the simple adjustment of an unnecessarily powerful definition. With the new definition implemented, Wiles' proof continues with no real alternations. This is not to trivialize McLarty's contribution. Whilst it's a simple matter to change a definition, it is certainly not a simple matter to demonstrate that Wiles' proof really can proceed untroubled by the new, weaker definition.

Mclarty (2020) has actually gone further and proven the more general result that the entirety of Grothendieck's main work, *Séminaire de Géométrie Algébrique du Bois Marie* (SGA)[4], can be founded in finite order arithmetic (a theory stronger than *PA* but weaker than *ZFC*). Again, as with Wiles' proof, McLarty shows that the proofs in SGA itself go through given a weaker notion of universes. He's showing that the original proof used unnecessarily strong axioms, not producing a new proof from a more conservative starting point.

Who proved Fermat's last theorem: Wiles or McLarty?
Wiles. Obviously.

[4] The SGA was published in eight volumes. To avoid gratuitous citations, a combined edition can be found online here: `http://library.msri.org/books/sga/`

McLarty's work here is exceptionally interesting, especially to those of us with more "philosophical" interests in mathematics but his discovery that Wiles was using universes, and his production of a work-around, doesn't take the proof away from Wiles. However, if one applies the bibliographical approach, it's McLarty's proof, not Wiles', which unpacks to a derivation from a commonly accepted foundation. The bibliographical approach seems to get the wrong answer to the question "Who proved Fermat's last theorem?"

Of course, one reply here could be that this case shows that Tarski-Grothendieck set theory is now just an acceptable foundation for one's proofs. That wouldn't be an entirely unreasonable claim. There's a strong case for large cardinal axioms. But if this was the case, one would have expected a little more conscious acknowledgement by the mathematical community, perhaps even some level of controversy. Whilst McLarty's results are widely known in proof-theory and related circles, the work has not received exceptionally wide-ranging attention beyond that.

What I think is a more plausible reading of the situation is that McLarty has shown how Wiles' actual proof can be unpacked into a derivation from more traditional foundations. Nothing major has changed in the best current version of the proof as a result of McLarty's contribution. What he's done is show how to shave down the present proof.

Returning to the programming analogy, I like to think of McLarty's contribution a little like a coder who's written a more efficient compiler. Wiles, the higher-level programmer, has written an exceptionally important algorithm which, using the compilers available at the time, needed a supercomputer to run. McLarty has written a compiler which can take the very same code but turn it into a much less computationally demanding ML script. McLarty's contribution is very helpful, but the algorithm is still Wiles'.

My explanation of this case, however, requires a rejection of the bibliographical approach to unpacking proofs. As with the cases of mistakes and multiple foundations, correctly unpacking a proof into a derivation requires some element of idealization. One can't just read the correct derivation from the bibliography.

So which derivation is *the* correct unpacking of Wiles' proof of Fermat's last theorem: The derivation via McLarty from finite order arithmetic or the derivation without his contribution from Tarski-Grothendieck set theory?

I don't believe either can be thought of as *the* correct unpacking. Both are perfectly good rational reconstructions of what's going on in the proof. But this would then be exactly what the pluralist is looking for: a proof which can be reasonably unpacked into multiple derivations with different axiomatic basies.

## 2.3 The Psychological & Counterfactual Approaches

An alternative approach to identifying a privileged derivation is a psychological approach. One might choose to privilege not necessarily what's "out there" but rather how the mathematician who wrote the proof actually thinks about their proof. Again, there's some prima facie plausibility to the claim that this particular unpacking should lead to a privileged derivation. After all, that derivation would be the derivation that the mathematician had in mind.

Two issues here:

First of all, many mathematics papers have multiple authors. There are two cases here: series and parallel.

In the parallel case, multiple authors work on the entire paper. In this case, perhaps those authors disagree about exactly which derivation is the one they had in mind. They might, for instance, disagree about how they'd want to represent higher-level mathematical concepts at lower levels. In this case, which derivation would be the preferred one?

In the series case, different mathematicians work on different sections of the paper. Again, if the mathematicians disagree on how to represent their proof at a lower level, it might not be possible to "daisy chain" together the derivations corresponding to different sections of the proof.

It's therefore unclear how to apply the psychological approach to multi-author papers whilst maintaining a single privileged derivation.

The second issue is that the psychological states of mathematicians likely aren't sufficiently detailed to uniquely determine a derivation.

The vast majority of mathematicians, especially those not working on foundational issues, do not know how to turn their proofs into derivations. It took Hales the best part of 20 years and a large research team to go from a proof of the Kepler conjecture to a full derivation. This was not merely a matter of getting down on paper what was already in his head (or perhaps in his notebooks, etc). It takes a lot of work and knowledge to unpack a proof into a derivation. Individual mathematicians likely haven't thought through the exact details of a derivation corresponding to their proofs. It's unlikely, then, that there's a single particular derivation that a mathematician might have in mind when producing a proof.

The psychological approach can potentially be adapted to avoid this concern. Perhaps what's worth considering is not the derivation that a mathematician actually had in mind but rather the one that they would produce if, like Hales, they decided to produce one.

This counterfactual approach would still encounter the problem of multiple authors: Different authors might give different derivations even under these ideal conditions. But the counterfactual approach at least avoids the (initial) issue of under-determination. It doesn't matter if a mathematician's actual psychological states under-determine which derivation is right, the counterfactual approach considers the facts of a world where they bothered to think through the details of a derivation, not the actual world.

That being said, the counterfactual approach might still lead to under-determination if there is not a single closest possible world where the mathematician completes a full derivation of their proof. It might be that the facts of the actual world are insufficient to determine exactly which of two possible derivations is the closest. It might be that the possible world where our hypothetical mathematician unpacks their proof into a derivation from *CTT+* is equidistant to the world where they do it in *ZFC*.

This is the known question of Stalnaker's Axiom, as discussed by Lewis (1973):
Is it always the case that either $X \mathrel{\Box\!\!\rightarrow} Y$ or $X \mathrel{\Box\!\!\rightarrow} \neg Y$?

Even if there are grounds to decide which world is close, a worry is that these are then irrelevant grounds for determining the most faithful unpacking of a proof. Suppose, for instance, that the reason why our mathematician would choose to use *ZFC* is something relatively benign. Perhaps they are slightly better friends with a Set Theorist than a Type Theorist and would choose on those grounds which foundations to use, so as to be able to collaborate with a friend. Those are good grounds to prefer one counterfactual scenario to another but strange grounds to prefer one unpacking of a proof over another.

As such, there's also good reason to be sceptical of the counterfactual approach.

## 2.4 In Summary

This section explains how the same proof can be unpacked, and I think unpacked reasonably and faithfully, into multiple different derivations, many of which would have different axioms. Three ways of privileging certain derivations over others are considered and dismissed. This is then, I hope, grounds not merely for claiming that proofs can be unpacked into derivations with different axiomatic basies, but that these different derivations are still honest, faithful unpackings of what's going on in the proof.

With the exception of proofs that have explicit axiomatic starting points, the relationship between ordinary proofs and axioms is clarified. The relationship is fuzzy. Proofs exist in a higher-level concept space without a singular axiomatic specification. It is rare that any particular axiomatic theory can be thought of as *the* axioms of a given proof. Instead, they should be thought of as a possible basis for an unpacking of the proof, but certainly not a unique basis.

## 3 Consequences for the Philosophy of Axiom Selection

Over recent years, a not inconsiderable amount of attention has been given to the question "Does Mathematics need new Axioms?"
This is discussed in Essay I: *On Axiom Selection*, but see Feferman et al. (1999) and the recent debate regarding the suitability of Category Theory as a foundation for mathematics[5].

Parties concerned seem happy to agree that (1) at present *ZFC* is *the* foundation for mathematics but disagree on whether (2) *ZFC* should continue to be the foundation for mathematics. Dissenters to (2) have or might propose extending *ZFC* (e.g. with the addition of some large cardinal axiom), reducing it (e.g. by moving to an intuitionistic logic or weakening choice), amending it (e.g. by replacing Choice with Determinacy, Maddy's the discussion of the Cabal seminar: Maddy (1988a) and Maddy (1988b)) or by taking an entirely non-set theoretic approach (e.g. the Univalent Foundations Program (2013)).

There has been, to my knowledge, no published disagreement with (1). There seems to be widespread acceptance at least within the published literature that *ZFC* is the actual present foundation of mathematics. Nevertheless, the conclusion of §2 casts some doubt on (1).

Exactly what does it take for some axiomatic theory to be the *actual* foundation for mathematics? Actual, as opposed to the foundation in some more idealized philosophical sense.

There are lots of ways an axiomatic theory might be the actual foundation for mathematics. Perhaps because it is widely regarded by mathematicians of all stripes as the foundation. Perhaps mathematicians could meet and take a vote! Perhaps the fact that an axiomatic theory frequently appears as the official foundation of prominent textbooks (e.g. many axiomatic set theory textbooks are explicit about their foundations). Perhaps there is a certain amount of delegation and that the relevant, specialist subdisciplines working on foundational mathematics come to one of the above kinds of agreement.

---

[5] SeeHellman (2003), Awodey (2004), Linnebo and Pettigrew (2011) and Ladyman and Presnell (2016), amongst others.

All that notwithstanding, one way that a particular axiomatic theory might be the actual foundation of mathematics is if it's the axiomatic basis of the majority of mathematical proofs. As discussed in §1, proofs themselves infrequently specify their axioms. Proofs need to be unpacked into derivations, which do explicitly state their axioms. But as §2 argues, most proofs can be (sufficiently faithfully) unpacked into derivations with different axiomatic basies. This means that for the majority of the mathematical literature, obviously excluding works in axiomatic mathematics that are explicit about their foundations, its possible to unpack them both to derivations from *ZFC* and to derivations from other axiomatic theories as well. But if that is true, then it's unclear that *ZFC* is *the* singular actual foundation for mathematics.

I suggest that the actual foundations of mathematics might be thought of in more pluralist terms. Instead of one foundation, perhaps there are many. Not a single axiomatic theory underneath it all but rather a cluster of distinct but often related axiomatic theories, connected to ordinary proofs by known definitions, tricks and methods, and a substantial amount of know-how by those concerned with foundational issues. When mathematicians need to do foundational work, they may pick any member of this cluster on the basis of their own preferences or practical needs safe in the knowledge that in all but the most marginal cases, they could have used any other.

This is, given the analysis of §2, at least the correct account of the actual foundations of the mathematical *literature*. Exactly what this means for the actual foundations of mathematics, as a whole, requires a discussion that I only hint at here.

Suffice for present purposes to say the following. My analysis of the relationship between proofs and axioms, if correct, provides some ground to be sceptical that there is a singular actual foundation for mathematics. Instead, a plural foundation is more plausible with mathematicians free to choose between a range of plausible and well-regarded axiomatic theories. I have no means demonstrated this completely. A great deal hangs on exactly what one means by the actual foundations of mathematics. But this analysis does provide some prima facie motivation for scepticism about a singular actual foundation for mathematics.

As a final point of comment, it should be noted that none of this speaks directly to the *normative* question of what the foundations of mathematics should be. Even if, in practice, there is no singular foundation for mathematics, that does not mean that there shouldn't be. That is a very different question, one addressed in Essay III: *Instrumentalism & Axiom Selection*.

That notwithstanding, there is a tendency in the philosophy of mathematics to be a little conservative. All else being equal, a philosophy of mathematics that is less disruptive of mathematical practice is better than one that is. If one adheres to that rule, then clearly my claims about the actual foundations of mathematics have significant philosophical upshot. They would shift the burden of proof onto the axiomatic monist. If the actual foundations of mathematics are plural, and granting this conservative approach, then it falls to the party wishing to adopt a singular foundation to explain why this change is warranted.

## 4 Conclusion

This paper builds on previous work on the relationship between mathematical proofs and derivations by considering the relationship between proofs and axioms. It argues that not only can proofs be unpacked to multiple derivations, but that this includes derivations from different axiomatic theories. In this sense, no singular axiomatic theory can be thought of *the* axioms of a proof, except in the case of proofs in axiomatic mathematics where the axioms are explicitly stated.

The reason for this result is explained. Mathematical proofs are written in a higher-level language, comparable to a higher-level computer language, which can be unpacked into lower-level formal languages in much the same way that higher-level computer code is compiled into machine code. But just as there are many possible computer compilers which return different ML script from the same code, there are different ways of unpacking proofs into more fundamental, formal languages. This pluralism is strong enough to even allow for different axiomatic theories to form the basis of derivations unpacked from the same proof.

Some initial consequences of this for the foundations of mathematics are outlined (though not entirely defended). This suggests, I believe, that mathematics actually has a plural foundation. There is not one singular axiomatic theory that is *the* basis for all of modern mathematics. Rather, there are many theories, all of which are sufficient to do foundational work. Mathematicians are free to choose whichever best suits their needs, when there is foundational work to be done.

## References


Andersen, L. E. (2020). Acceptable gaps in mathematical proofs. *Synthese*, 197(1):233–247.

Awodey, S. (2004). An answer to hellman's question: Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 12(1):54–64.

Azzouni, J. (2004). The derivation-indicator view of mathematical practice. *Philosophia Mathematica*, 12(2):81–106.

Burgess, J. and Toffoli, S. D. (2022). What is mathematical rigor? *Aphex*, 25:1–17.

Cohen, P. J. (1963). The independence of the continuum hypothesis. *Proceedings of the National Academy of Sciences of the United States of America*, 50(6):1143–1148.

Detlefsen, M. (2008). Proof: Its nature and significance. In Gold, B. and Simons, R. A., editors, *Proof and Other Dilemmas: Mathematics and Philosophy*, page 1. Mathematical Association of America.

Fallis, D. (2003). Intentional gaps in mathematical proofs. *Synthese*, 134(1-2):45–69.

Feferman, S., Friedman, H. M., Maddy, P., and Steel, J. R. (1999). Does mathematics need new axioms. *Bulletin of Symbolic Logic*, 6(4):401–446.

Goldfarb, W. D. (1984). The unsolvability of the gödel class with identity. *The Journal of Symbolic Logic*, 49(4):1237–1252.

Gödel, K. (1933). Zum entscheidungsproblem des logischen funktionenkalküls. *Monatshefte für Mathematik und Physik*, 40:433–443.

Hales, T., Adams, M., Bauer, G., Dang, T. D., Harrison, J., Hoang, L. T., Kaliszyk, C., Magron, V., McLaughlin, S., Nguyen, T. T., Nguyen, Q. T., Nipkow, T., Obua, S., Pleso, J., Rute, J., Solovyev, A., Ta, T. H. A., Tran, N. T., Trieu, T. D., Urban, J., Vu, K., and Zumkeller, R. (2017). A formal proof of the kepler conjecture. *Forum of Mathematics, Pi*, 5:e2.

Hales, T. C. (2005). A proof of the kepler conjecture. *Annals of Mathematics*, 162(3):1065–1185.

Hamami, Y. (2014). Mathematical rigor, proof gap and the validity of mathematical inference. *Philosophia Scientiae*, 18(1):7–26.

Hamami, Y. (2022). Mathematical rigor and proof. *Review of Symbolic Logic*, 15(2):409–449.

Hamami, Y. and Morris, R. (forthcoming). Rationality in mathematical proofs. *Australasian Journal of Philosophy*.

Hellman, G. (2003). Does category theory provide a framework for mathematical structuralism? *Philosophia Mathematica*, 11(2):129–157.

Kitcher, P. (1981). Mathematical rigor–who needs it? *Noûs*, 15(4):469–493.

Ladyman, J. and Presnell, S. (2016). Does homotopy type theory provide a foundation for mathematics? *British Journal for the Philosophy of Science*, page axw006.

Leitgeb, H. (2009). On formal and informal provability. In Bueno, O. L. O., editor, *New Waves in Philosophy of Mathematics*, pages 263–299.

Lewis, D. K. (1973). *Counterfactuals*. Malden, Mass.: Blackwell.

Linnebo, O. and Pettigrew, R. (2011). Category theory as an autonomous foundation. *Philosophia Mathematica*, 19(3):227–254.

Maddy, P. (1988a). Believing the axioms i. *The Journal of Symbolic Logic*, 53(2):481–511.

Maddy, P. (1988b). Believing the axioms ii. *The Journal of Symbolic Logic*, 53(3):736–764.

Marfori, M. A. (2010). Informal proofs and mathematical rigour. *Studia Logica*, 96(2):261–272.

McLarty, C. (2010). What does it take to prove fermat's last theorem? grothendieck and the logic of number theory. *Bulletin of Symbolic Logic*, 16(3):359–377.

Mclarty, C. (2020). The large structures of grothendieck founded on finite-order arithmetic. *Review of Symbolic Logic*, 13(2):296–325.

Tanswell, F. (2015). A problem with the dependence of informal proofs on formal proofs. *Philosophia Mathematica*, 23(3):295–310.

Tanswell, F. S. (2016). Proof, rigour and informality : A virtue account of mathematical knowledge. In *St Andrews Research Repository Philosophy Dissertations*.

Toffoli, S. D. (2021). Groundwork for a fallibilist account of mathematics. *Philosophical Quarterly*, 7(4):823–844.

Toffoli, S. D. (forthcoming). Who?s afraid of mathematical diagrams? *Philosophers' Imprint*.

Univalent Foundations Program, T. (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. `https://homotopytypetheory.org/book`, Institute for Advanced Study.

Wiles, A. (1995). Modular elliptic curves and fermat's last theorem. *Annals of Mathematics*, 141(3):443–55.

# Essay V: Langauge, Truths & Logics:

## A Defence of (Neo) Carnapian Logical Pluralism

## Introduction

Carnap's logical pluralism is a well know, though not highly regarded, version of logical pluralism. The view is mentioned in almost all extended discussions of logical pluralism but is quickly dismissed as either (1) the victim of one or multiple devastating objections or (2) obviously true but entirely uninteresting and besides the point.

In this paper, I argue that Carnapian logical pluralism survives all the major objections against it. Moreover, it is an interesting type of pluralism that is relevant to important questions about logic.

At certain points, the view I defend is not necessarily exactly Carnap's view. If Carnap's actual view survives all the objections put to it rests on some exegetical questions that I won't attempt to resolve here. The view that I defend is very much in the spirit of Carnap's logical pluralism, but (potentially) breaks from it at certain critical moments. Crucially, my view abandons Carnap's tolerance principle. I distinguish between a logic being descriptively correct for some language and its normative status as a viable logic for actual reasoning. I refer to my view (without Tolerance) as the Neo-Carnapian view and Carnap's view (with Tolerance) as the classical Carnapian view and refer to both of them together as the Carnapian views.

§1 and §2 set out the groundwork. §1 clarifies precisely which kind of pluralism the Carnapian views are committed to. §2 then explains both the Carnapian views, discussing their differences in §2.2.

§3, by far the most substantial, discusses what I take to be the four main objections to Carnapian Logical Pluralism, defending the Neo-Carnapian position, and sometimes the classical Carnapian position to boot. The objections considered are as follows. The "Wrong type of Pluralism" objection (§3.1), which claims that the (Neo) Carnapian logical pluralism is an uninteresting, irrelevant or insubstantial version of logical pluralism (Cook, 2010). Prior (1960)'s and Belnap (1962)'s Tonk objection (§3.2). Next, I consider the meaning objection, originally from Quine (1970) but later re-introduced into the modern debate by Restall (2002) and discussed in Griffiths and Paseau (2022) (§3.3). Lastly (§3.4), a class of objection, which I call the Meta-Logic Objections, which includes Quine (1936)'s argument against logical truth by convention (§3.4.1) and what I call the mismatch in Griffiths and Paseau (2022), amongst other places (§3.4.2).

# 1 The Carnapian Approaches to Logical Correctness

## 1.1 Logical Pluralism via Alethic Pluralism via Semantic Pluralism

Both the Neo and classical Carnapian views are examples of inter-linguistic logical pluralism. They hold that there are many correct logics because there are many logics correct for some language.

A logic is correct for a language L iff it is sound and complete[1] with respect to truth-in-L. Conversely, I say that a language L *satisfies* a logic iff that logic is correct for L. However, different languages have different semantic features with give rise to different notions of truth.

[1]An exception will need to be made for higher-order logics which might not admit complete logics. In this case, some condition short of completeness should be applied.

Consider the following example.

In classical logic a predicate's extension (the things that make it true) and anti-extension (the things that make it false) are both (1) disjoint and (2) jointly exhaustive of the domain[2]. In conjunction with the fact that all names refer, this means that all atomic sentences are either true or false, but never both. In familiar ways, connectives can be defined truth functionally and the familiar inference rules of classical logic can be shown to be valid.

However, the exclusion and exhaustion of the extension and anti-extension is not a necessary semantic fact. It's possible to build languages where either (1) the extension and anti-extension can overlap but are jointly exhaustive (2) they cannot overlap but are not jointly exhaustive or (3) they neither need to overlap nor be jointly exhaustive. (1) leads to truth-value gluts, (2) to truth-value gaps and (3) to both. Again, truth-functional connectives can be defined. Familiar non-classical inference rules will be true of these connectives.

Both Carnapian views then tell a similar, simple story. Logics are true of a language iff they are sound and complete for truth in that language. But the nature of truth in a language depends upon linguistically contingent facts, such as the rules governing of admissible extensions and anti-extensions of predicates. So the correct logic is a linguistically contingent matter. There are different logics correct for some language.

For both Carnapian views, there is nothing more to say about the correctness of a logic than this. There's no language-independent notion of truth or constraint on logical correctness beyond correctness for some language.

This is important as it distinguishes the Carnapian claim from the functionally trivial claim that there are many logics, in the sense of many definable mathematical structures which satisfy whatever structural constraints it takes to be a logic. They're making a substantive and contested claim about what facts are relevant for deciding if a logic is *correct*.

[2]This is a slightly atypical presentation of classical logic. Orthodox presentations typically don't differentiate between an extension and anti-extension in classical logic, as the anti-extension is just the complement of the extension in the relevant $n^{th}$ Cartesian product of the domain. However, the atypical presentation is useful for present comparison.

As a final remark here, I don't intend the Neo-Carnapian view to be restricted to artificial languages. I include natural languages too. This is consequential in response to Quine's Conventions argument (§3.4.1). I don't know if Carnap intended to restrict his inquiry to artificial languages or not. If he did, then this is a point of difference between the Neo and classical Carnapian views.

What makes a semantic theory the correct semantic theory for a language?
At the least, it should play the appropriate role in a descriptively accurate theory of speaker utterances and inferences. But such a theory would also need an account of pragmatics, conversational norms, the psychology of speakers, speaker beliefs and much more. Linguistics is complicated and the role of semantics in predicting speaker behaviour is somewhat abstracted, but certainly present.

From a correct semantic theory, which will undoubtedly be more baroque than its formal cousins, one can extract the (or perhaps merely a) correct theory of truth for a natural language. The (or a) correct logic of this natural language will then be the one sound and complete for truth in that language.

Now, I'm very sympathetic to Shapiro's version of logical pluralism arising here. There might be multiple different semantic theories which can play the appropriate role in a descriptively accurate linguistic theory. This is in no small part because speaker behaviour might radically underdetermine key data points. I doubt there's all that much uniformity, for instance, on how actual speakers deal with the Liar paradox, though this is armchair linguistics.

But if there are many semantic theories correct for some language, then there can be many theories of truth correct for some language and consequently many logics. This is a very carnapianized version of Shapiro (2014)'s logical pluralism. I think it's plausible, but I won't defend it any further here. It's worth noting that this is another point of difference between me and Carnap, who rejected this claim.

### 1.2 Carnap's Tolerance Principle

One of the most notable, and so far unmentioned, parts of Carnap's philosophy of logic[3] is his tolerance principle. In Carnap LSL he states the following:

> *"In logic there are no morals. Everyone is at liberty to build up [their] own logic, i.e. [their] own language. All that is required of [them] is that, if [they] wishes to discuss it, [they] must state [their] methods clearly, and give syntactical rules instead of philosophical arguments."* (Carnap, 1934, p52)

He continues saying that: *"It is not [the business of philosophy of logic] to set up prohibitions, but to arrive at conventions"*.

There are two ways of reading this, one quite far from the Neo-Carnapian position and one quite close to it.

On the first reading, Carnap is making a claim about the kinds of logic one is entitled to use when reasoning. He's claiming that the only restriction on the kinds of logic one may use when reasoning is that it's true in some language. Presumably, also, that one adopts this language. On this reading, there are no grounds on which to choose between languages. This is just a matter of practical convention setting. As he says, in logic there are no morals.

An analogy here would be between deciding which side of the road to drive on when designing traffic laws. There's no objectively correct answer to this. Road systems function perfectly well in both cases. There aren't even any serious normative considerations in favour of one choice or the other. It's simply a matter of settling which of a set of equally good options we're going to go forwards with.

Call this the non-normative reading of Carnap.

[3] This paper discusses the *Logical Syntax* (Carnap, 1934) and *Empiricism, Semantics and Ontology* (Carnap, 1950), but see also Carnap's *Aufbau* (Carnap, 1928) and *Meaning and Necessity* (Carnap, 1947).

In contrast, the second reading puts a bit more normativity into Carnap. In his later work, Carnap (1950) discusses choosing between nominalist and non-nominalist languages. In the closing lines of *Empiricism, Semantics and Ontology* he says the following:

*"The acceptance or rejection of abstract linguistic forms, just as the acceptance or rejection of any other linguistic forms in any branch of science, will finally be decided by their efficiency as instruments, the ratio of the results achieved to the amount and complexity of the efforts required."*

Here Carnap clearly is applying normative criteria to language (and hence logic) selection. Logic has a function, a purpose. This is not to be understood metaphysically but very practically. Logics can be decided upon, though perhaps not up to uniqueness, on practical, normative grounds. Whilst there are no *descriptively* correct or incorrect logics, provided one can find a language which satisfies them, there are *normatively* better or worse logics. Contra the first Carnap, when it comes to use, it's not simply that any logic goes.

Call this the normative reading.

This second Carnap, the normative Carnap, is much closer to the Neo-Carnapian. The Neo-Carnapian takes the normative reading and builds on it.

The Neo-Carnapian differentiates between the question of whether a logic is correct (for some language) and whether they should use it. That a logic is correct for some language might as much be a reason to not use that language as it is permission to use the logic.

The Neo-Carnapian is unashamedly instrumental in their approach to language selection. They see logic as a tool and will pick whichever language gives them the most useful tool. They see a web of reasons, both practical and theoretical, for using particular logics in particular contexts and for particular tasks. For the Neo-Carnapian, this is the crux of the philosophy of logic: the very practical question of the benefits and drawbacks of using this logic or that.

For the Neo-Carnapian, in logic there are only morals.

## 2 Pluralism Classified

For any given subject, there are often many different kinds of pluralism one might adopt. A pluralism of pluralisms. When being exact, it's not enough to simply say "logical pluralism" or "alethic pluralism" or suchlike. There are typically many different independent positions all of which can reasonably go by the name "X Pluralism".

With respect to logical pluralism, the following are all positions that could go by the name "logical pluralism":

- There are many logics
- There are many correct logics
- There are many correct theories of inference
- There are many logics correct for some language
- For some languages, there are many correct logics
- There are many logics with which one may reason

Whilst these different notions of logical pluralism are certainly interconnected and *typically* if one is a pluralist about some of this, one will be a pluralist about many more, they should be understood as separate views which can come apart, at least in principle. This is similarly so for monism and nihilism.

This section clarifies in exactly which ways the Carnapian views are pluralist.

As a more general point, I hope this section also shows that these kinds of subtle distinctions raised above are highly consequential. They do come apart and they come apart in important moments. On the basis of this, I hope that we stop thinking of logical pluralism (and monism, and nihilism) as a singular position, but rather a family of connected but strictly independent views.

## 2.1 The Scope Question

As mentioned above, on the Carnapian views logics are only ever correct with respect to a particular language, in much the same way that things like theories of grammar are only correct with respect to a particular language. The question "How many theories of grammar are correct?" is ambiguous and only makes sense upon clarification. One might ask:

1. Wide scope question: How many grammars are correct for all languages?

2. Medium scope question: How many grammars are correct for some language?

3. Narrow scope question: For a given language, how many grammars are correct?

4. Objective question: How many grammars are objectively correct?

And for each of these questions one might answer "none" and be a nihilist, "exactly one" and be a monist or "many" and be a pluralist.

The question "How many logics are correct?" is ambiguous in the same way. One might ask:

1. Wide scope question: How many logics are correct for all languages? (None)

2. Medium scope question: How many logics are correct for some language? (Many)

3. Narrow scope question: For a given language, how many logics are correct? (Me: many, Carnap: one)

4. Objective question: How many logics are objectively correct? (Carnap: None, Me: see other work)

Both the classical and Neo-Carnapian answers are given in the brackets above.

Note that this scope distinction holds for other related kinds of pluralism too. The question "How many correct alethic/semantic theories are there?" is ambiguous and could mean:

1. Wide scope question: How many alethic/semantic theories are correct for all languages?
2. Medium scope question: How many alethic/semantic theories are correct for some language?
3. Narrow scope question: For a given language, how many alethic/semantic theories are correct?
4. Objection question: How many alethic/semantic theories are objectively correct?

In fact, the Carnapian views defend medium-scope logical pluralism on the basis of medium-scope alethic pluralism, which is in turn defended on the basis of medium-scope semantic pluralism.

Now I don't want to claim that any one of these precisifications of the initial question "How many logics are there?" is the *correct* version of the question, or even the best or most interesting. I resist efforts by, e.g., Cooke to narrow the question of logical pluralism simply to the narrow scope question. I see logical pluralism not as a singular position, but as a family of interconnected but strictly distinct positions, united by the fact that they give a pluralist answer to some reasonable precisification of the question "How many logics are there?".

Both the classical and Neo-Carnapian views defend medium-scope logical pluralism. They hold that logical pluralism is true because there are many logics correct. This actually puts them apart from most contemporary versions of logical pluralism, which typically focus on the narrow-scoped question. Both Beall and Restall's, and Shapiro's versions of logical pluralism are narrow-scoped, for instance.

## 2.2 Correctness Pluralism vs Validity Pluralism

Another important distinction to be made is between correctness pluralism and validity pluralism. This is important when contrasting the Carnapian views with Beall and Restall (2005)'s, and also for the reply to the Meaning Objection (§3.3).

The correctness pluralist claims that there are many correct logics. Perhaps many correct for all languages, for some language, for a given language, perhaps objectively. The Carnapian views are correctness pluralists.

The validity pluralist claims that there are many correct theories of valid inference. In other words, there are many correct theories of which sentences follow from which others.

These might, at first sight, appear to come together. Perhaps validity pluralism is true iff correctness pluralism is true. After all, logics are theories of valid inference.

The claims come apart, however, because information about which language one is working in is encoded in sentences themselves. Consider the validity of the following informal arguments:

1. All men are mortal. Socrates is a man. So Socrates is mortal.
2. Alle Menschen sind sterblich. Sokrates ist ein Mensch. Also ist Sokrates sterblich.

As should be obvious, the lexicon of the sentences restricts, and in these cases determines, the language in which the arguments are to be evaluated.

In the case of German and English, this is inconsequential to the evaluation of the argument as the two languages share a logic. In the case of languages satisfying different logics, the difference is consequential. Consider the following case. There are two formal languages. One is classical. The other is a language with truth value gluts and the connectives of the logic of paradox. The languages are unambiguously delineated in their lexicon. In particular, the subscripts *CL* and *LP* are applied to the traditional symbols for the connectives to indicate if the connective belongs to the first or the second language. So $\rightarrow_{CL}$ is the conditional from the first language and $\rightarrow_{LP}$ the conditional from the second.

Consider the following two arguments:

1. $\phi \vee_{CL} \psi,\ \neg_{CL}\phi \vdash \psi$
2. $\phi \vee_{LP} \psi,\ \neg_{LP}\phi \vdash \psi$

Despite their prima facie similarity, the first statement is true whilst the second is false. Disjunctive syllogism is invalid in *LP*. But here we can see that the connectives tell us which language the sentences come from. This tells us the underlying theory of truth that should be used to determine the validity of the argument, and consequently which logic.

Returning to the difference between validity and correctness pluralism, only one theory of validity is needed to capture this case. This is the hybrid theory which says from classical statements use classical logic to infer classical conclusions, but from *LP* statements use *LP* inferences to infer *LP* conclusions. For the (medium scope) correctness pluralist, there are two correct logics. Classical logic is correct for the first language and *LP* is correct for the second.

In this case, validity pluralism is false but correctness pluralism is true.

Now, again, I don't wish to say that either correctness pluralism or validity pluralism is *the* correct notion of logical pluralism. As stated above, logical pluralism is a family of views not a singular view. Beall and Restall's pluralism is a kind of validity pluralism. They believe that there are multiple interpretations of the consequence relation, with different logics correct given different interpretations. Their view is certainly a plausible approach to logical pluralism, though one I do express some scepticism towards[4]. Validity pluralism is one amongst many ways one might be a pluralist. But it's certainly not the Carnapian approach to pluralism.

[4]In particular I'm sceptical that: (1) The widest scope consequence relation can't simply be treated as the correct consequence relation and (2) Even if our actual notion of consequence in, say, English is ambiguous between different interpretations, this might just be a reason to conceptually engineer a more precise consequence relation, rather than sit with the ambiguity. They need to show me the *value* of keeping multiple consequence relations in the language, rather than adapting our language to clarify the ambiguity. Both of these are points on which I could be convinced otherwise.

## 2.3 Descriptive Pluralism vs Normative Pluralism

The final distinction to be made relates to the tolerance principle (§1.2) as well as to the Tonk objection (§3.2). The Neo-Carnapian stresses a distinction between the descriptive question of if a logic is correct (for some language) and the practical, normative question of if one may therefore reason with it.

A disagreement already exists within the literature as to whether logics are normative theories of reasoning or descriptive theories of truth (See Blake-Turner and Russell (2018) and Russell (2020) for discussion). The disagreement, I think, is misplaced. Logics are not inherently either. Logics are simply formal systems. Logics can, however, be *used* as normative theories of reasoning or descriptive theories of truth. Typically the former in virtue of the latter, as Russell notes. But this is not a property of logics themselves, it's a feature of how we use them.

The distinction between the two comes apart in cases of poorly designed languages. Consider, for instance, a classical language but without a complete set of connectives. Suppose, for instance, the language only has conjunction, disjunction and a conditional but not negation. There will be some sub-classical logic correct for this language. Should one reason with this logic?

No. There's no practical benefit to removing negation from the system. It would clearly be much better to add it back in and move to classical logic proper. Nevertheless, one can't deny that the logic is correct for some language, just not a very useful language.

The distinction between descriptive and normative correctness should then be thought of like this, at least within the context of the Carnapian's medium-scoped pluralism[5]. There are a number of languages. Each of these languages satisfies at least one logic. However, these languages are not normatively equal. They can be evaluated and compared along a number of possible dimensions.

[5] The distinction can easily be adapted to other versions of logical pluralism. Take Shapiro (2014)'s model pluralism. Even if there are many descriptively correct models, these might not all be normatively on a par. Some might be better models than others. The challenge is, I think, quite cutting against Beall and Restall's logical pluralism. Even if there are many entailment relations, it seems very plausible that only one of these will be relevant to the norms of reasoning.

What it means for a logic to be normatively correct is for it to be descriptively correct for some language which is appropriately normatively endorsed. Exactly what "appropriately normatively endorsed" means will depend on the exact normative structure applied to the languages. But to give the simple binary case, the logic must be descriptively correct for a language with which one may reason.

Normative logical monism, pluralism and nihilism are then simply monism, pluralism or nihilism about normative logical correctness.

Within this paper, I largely discuss descriptive, rather than normative, correctness. An exception to this is §3.2 which discusses the Tonk objection. Here both are discussed.

The classical Carnapian is definitely a normative logical pluralist. The tolerance principle means that there are many logics with which one may reason.

The Neo-Carnapian view is compatible with normative logical monism. Nothing said in this paper rules that position out. That being said, I am a normative logical pluralist. I believe there are many useful logics with which one may reason, but perhaps not that many. I also believe that's the more natural position to take as a Neo-Carnapian, but more because it's in keeping with the general attitude of the view, not because of any strict argument connecting one to the other.

All that notwithstanding, the type of pluralism that the Neo-Carnapian commits to is descriptive, not normative, pluralism.

# 3 Objections

## 3.1 Objection 1: The Wrong Type of Pluralism?

In his well-known survey of logical pluralism, Cook (2010) draws a distinction between logical pluralism, in general, and what he calls "substantial" logical pluralism. Whilst Cook agrees with my analysis in §2 that there are many different views that might go by the name "logical pluralism" Cook takes one of these questions in particular to be privileged, the "substantial" version of logical pluralism.

For Cook, substantial logical pluralism is a kind of narrow-scope descriptive correctness pluralism where the fixed language in question is some natural language.

The classical Carnapian view is then dismissed as insubstantial as its pluralism is medium rather than narrow scoped. Note, the Neo-Carnapian view does not differ from the classical view in this regard so, if Cook is right, both Carnapian views are at risk.

Cook is not all that explicit about why narrow-scoped logical pluralism is substantial but medium-scoped logical pluralism is not. He simply applies the term "substantial" to the narrow-scoped version without all that much in the way of justification.

Now cook is writing a survey article, not an argumentative piece. It's perhaps unfair to expect too much in the way of justification, especially if it were to get in the way of exposition. However, I think there are two arguments which can be extracted from the paper.

The first of these is that Carnapian logical pluralism should really be seen as relativism, not pluralism. As the worry goes, the Carnapian does not think that there are many correct logics. For any one reasoner at any one time, there is only one correct logic, but this is a function of the language they are using. Logical correctness is then singular but language relative.

If this is Cook's objection, it confuses the type of logical pluralism that the Carnapian commits to. All that analysis amounts to is the claim that the Carnapian is not a narrow-scoped monist. But the Carnapian never claims to be, their pluralism is medium-scoped. One can be a monist about one version of pluralism but a pluralist about another.

The second argument that Cook might have been making is that Carnapian logical pluralism is insufficiently controversial; it gets logical pluralism "on the cheap", to use his phrase. Cook is keen that logical pluralism be a controversial thesis. Cook might be concerned that the Carnapian view is an overly easy or cheap version of pluralism.

Three replies to this.

First, historically, the Carnapian view has been taken to be all but refuted. See, for instance, commentary by Soames (2003), Sider (2011) and Benacerraf (1973), discussed below in §3.3. If it's controversy Cook is searching for, defending a Carnapian view against half a century of orthodoxy would certainly achieve that!

Second, as I stress in §1.1, the Carnapian views are committed to more than just the claim that there are many languages which satisfy different logics. That claim would not amount to much more than what Cook calls mathematical logical pluralism, the claim that there are many logics, in a purely mathematical sense. The Carnapian views commit to the more substantive claim that there's no more to logical correctness than correctness for some language. That is a substantive claim.

Consider, for instance, the debate between Logical Realists like Tahko (2021) and McSweeney (2019) who think there are objective language-independent logical facts. Tahko, in particular, acknowledges that there are many languages which satisfy different logics, he just thinks that the relevant notion of logical correctness is correctness with respect to the objective logical facts, not correctness with respect to the semantic facts.

The Carnapian and the Logical Realist, therefore, agree on the claim that there are many languages which satisfy many logics. What the logical realist demands of a logic in addition to this is that the logic is correct with respect to some objective logical facts, whatever those might be.

The Carnapian's claim that all there is to logical correctness is correctness for some language is by no means a trivial part of their view.

See Essay VI: *Three Approaches to Logical Correctness* for more discussion.

Finally, controversy or contentiousness is not an especially good measure of philosophical consequence. As an example it is, I take it, relatively uncontroversial to claim that Gettier (1963) problems refuted the classical tripartite conception of knowledge. Nevertheless, this is (and was) an immensely consequential fact. So even if medium-scoped Pluralism turns out to be uncontroversial, that doesn't mean it's not highly consequential.

Continuing this line of thought, and closing out this section, I want to meet Cook's challenge positively and explain why I think we should care about medium-scoped pluralism. In other words, why the Carnapian views are consequential.

Now I'm not claiming this to the exclusion of the narrow-scoped question, which I also take to be interesting and consequential. As I think the latter part of Cook's paper shows, there are some immensely interesting narrow-scoped views that have been outlined in recent years. I'm just objecting to what I take to be an unwarranted narrowing of the debate to *only* considering narrow-scoped pluralism.

Why does Carnapian logical pluralism matter?

Carnapian logical pluralism matters for the very practical and important question of what logics we may use when reasoning. When reasoning we are not locked into a single language. We can choose to revise our language as part of a debate, discussion or chain of reasoning. When faced with an unclear or ambiguous concept, we might clarify it. When faced with a concept that is dysfunctional, we might choose to re-engineer it all together. This is the idea behind conceptual engineering as an approach to philosophy. In philosophy, our concepts are not fixed, they can be revised on the basis of a whole host of considerations. The ultimate extension of that is what you could call "logic engineering". Just as we can revise more local parts of our language, we choose to re-engineer the entire structure of our language, perhaps to the extent that we have moved into an entirely new language altogether.

For instance, suppose we discover, as is likely the case in most natural languages, that our language contains vague predicates. Suppose that for the purposes at hand a fuzzy logic is not helpful; we have reasons to prefer classical logic in this context. It would be very sensible in this context to change our language by applying some procedure for collapsing the vague predicates into Boolean ones, e.g. applying arbitrary cut-offs to the predicates. But note then that our linguistic re-engineering results in a change of logic. Within the scope of conceptual engineering, there's nothing special or sacred about logical rules. They can be changed as readily as any other part of language.

But this note that a crucial observation in order for this kind of procedure to be valid is the observation that logical correctness is really only a matter of which language one chooses to work in. It is only possible to "logic engineer" in this way if one accepts (1) that logical truths are semantic and that (2) logical rules are semantically contingent, they can be changed by changing the language.

Now this leaves out the most interesting question. Just because we can re-engineer our logic during an inquiry, doesn't mean that we should. When should we, if at all? Is there one logic that is the best one to use in all circumstances or many (normative logical pluralism)? On these questions, nothing is presently said. But for these questions to even make sense in the way outlined here, one has to first adopt Carnapian logical pluralism.

But note, returning to Cook (2010)'s objection, that it's medium-scope pluralism that's relevant to these questions, not narrow-scope pluralism. If these questions are worth answering, Cook's narrowing of the debate is unjustified.

## 3.2 Objection 2: The Tonk Objection

### 3.2.1 The Objection Explained

A famous objection to Carnapian logical pluralism is Prior (1960)'s Tonk objection.

Prior reads Carnap as claiming that any definable logic, in the sense of any clearly stateable set of inference rules, may be used in reasoning, providing that one is clear about the rules one uses.

Looking at Carnap's "In logic there are no morals" quote in isolation, it's understandable to see why prior things this. Carnap says the following:

*"In logic there are no morals. Everyone is at liberty to build up [their] own logic, i.e. [their] own language. All that is required of [them] is that, if [they] wishes to discuss it, [they] must state [their] methods clearly, and give syntactical rules instead of philosophical arguments."* LSL P52

If by "syntactical rules" one takes Carnap to mean inference rules, then Prior's reading of Carnap is accurate. All that is required is to be allowed to use a logic is to clearly state the inference rules.

In response to this, Prior defines the so-called nonsense connective "Tonk". Tonk is a binary connective with the classical introduction rule of or and the classical elimination rule of and. Let $\star$ be the tonk connective.

Formally: $\phi \vdash (\phi \star \psi)$ and $(\phi \star \psi) \vdash \psi$

But logics containing Tonk are (typically) trivial.
Let $\top$ be some theorem. Then:
$\top \vdash (\top \star \phi)$
$(\top \star \phi) \vdash \phi$
So together: $\top \vdash \phi$.

If a logic containing Tonk has even one theorem, then all formulae are theorems and the logic is trivial. This is, according to Prior though I agree, not an acceptable consequence. One may not reason with Tonk-logics.

One possible reply came quickly from Belnap (1962). A simple adjustment to the Carnapian position, in fact, one in line with how I present Carnap's view, is to substitute syntactic rules for semantic rules within the principle. Instead of allowing any syntactically definable logic to be used, instead one may only use a logic if one can define a semantic structure (i.e. a language) in which the rules are sound.

This is just my non-normative reading of Carnap. Despite the above quote, I think there's good reason to think that this is what Carnap had in mind all along. In the logical syntax, Carnap doesn't yet have a fully formed notion of the syntax-semantics distinction. Kouri (2019) even calls this reading the "standard interpretation". Bellnap's reply to Prior was really just Carnap's view all along.

All of this is a little beside the point, though. One can define languages which satisfy Tonk logics. Take, for instance, a language with only one truth value, which is designated. If every sentence is true in some language, then tonk logics would be sound and complete for this language. So Bellnap's reply doesn't escape the objection.

#### 3.2.2 The Reply

The Neo-Carnapian reply (also available to my normative reading of Carnap) is to accept that Tonk logics are the correct logics for certain languages but to reject the inference to the normative claim that one may therefore reason with Tonk. As outlined above (§1.2 and §2.3), the Neo-Carnapian makes a distinction between the descriptive correctness of a logic (i.e. the correctness of a logic for some language).

That a language satisfies some logic does not mean that one may use that logic in reasoning. It is not a consideration in favour of a logic that it is satisfied by some language. Instead, that some language satisfies some undesirable logic is simply a reason not to adopt the language.

Because they are trivial, tonk logics are not practically useful logic. There is, I suspect, no situation in which reasoning with Tonk would be helpful. Consequently, the discovery that some languages satisfy Tonk logics is simply reason to avoid those languages.

The Neo-Carnapian, therefore, accepts the consequences of the Tonk argument on the descriptive level but rejects it on the normative level. Tonk logics are correct for some languages, but that does not mean we can reason with them.

### 3.3 Objection 3: The Meaning Objection

A very popular objection to Carnapian logical pluralism is the meaning objection. It originates in Quine (1970) but can be found more recently in a paper by Greg Restall (2002) and as one of the two replies to Carnapian logical pluralism given in Griffiths and Paseau (2022).

In brief, the objection claims that Carnapian logical pluralism is only based on ambiguity in the language expressing an inference, not on real pluralism about the validity of the rules themselves.

The argument proceeds as follows.

Carnapian logical pluralism relies on variation of the correctness of logical rules between languages. Disjunctive syllogism, for instance, is valid in classical languages but invalid in languages with truth-value gluts. Given any sentence, if one knows to which language it belongs, it is always clear which other statements follow from it and which do not[6].

But, so contends the objection, either a sentence is ambiguous or we know to which language it belongs. Suppose that there are some sentences that exist in two languages but with different implications in each. For instance $(\phi \vee \psi)$ and $\neg\phi$ might exist in both a classical and a glutty language, but only entail $\psi$ in the former. The implications of these sentences are different because they have different truth conditions in these languages. But if they have different truth conditions they are, in fact, not the same statements. They are different statements expressed using the same signs. The apparent pluralism is really just a result of an ambiguity relating to the signs used[7]. As soon as the meanings of $\vee$ and $\neg$ are clarified, the apparent pluralism falls away.

[6]As a side note, the Neo-Carnapian is comfortable combining their view with various kinds of narrow-scope pluralism so might not grant this. But, if this argument is successful, it then would be this additional commitment which is doing all the important work.

[7]As a side note, It's interesting that Greg Restall has made this argument against Carnapian logical pluralism, given that his entire account of logical pluralism rests on the purported ambiguity of the $\vdash$ sign. If pluralism arising from ambiguity is problematic, Restall's own position falls foul of an analogue objection. As soon as the meaning of "entails" is clarified, any apparent pluralism falls away.

This connects back directly to the discussion in §2.1 on validity and correctness pluralism. The sentences in some entailment relation contain information about which language they belong to. The meaning objection simply takes this further arguing that they always contain enough information to completely determine the language and hence the validity of the inference.

There are two key claims in this objection worth highlighting:

1. For a view to be pluralist, it needs to be pluralist about which inferences are valid.
2. If an inference is valid in one language but invalid in another it is a result of different truth conditions in some of the premises.

Both of these claims can be challenged.

First of all, the meaning objection simply misunderstands the kind of pluralism that the Carnapian (Neo or classical) is arguing for. They don't claim to be a validity pluralist. They're a correctness pluralist. It wouldn't undermine the Carnapian's position in the slightest to accept that any apparent instances of validity pluralism are really just instances of linguistic ambiguity. It would still be the case that there are many different languages that satisfy different logics.

On the second point, there are many instances of differences in the validity of the logical rules that do not result in a difference in the truth conditions of the premises involved. Two examples.

First, consider a language which allows for models with empty domains, but which is otherwise classical. Compare this to a classical language. Consider the statement $\forall x\phi(x) \vdash \exists x\phi(x)$. This is invalid in the first language but valid in the second, modulo certain assumptions about how to treat quantification over empty domains[8].

[8] I'm assuming that any universal statement over an empty domain is true whilst the corresponding existential statements will be false. Other ways of handling quantification over empty domains will result in other differences in validity. For instance if universal statements get truth value gaps in empty models, then $\neg\exists x\phi(x) \vdash \forall x\neg\phi(x)$ is invalid in the language with empty models but valid in the classical language.

Here, the difference in the validity of the rule is not a result of the meanings of the terms involved. The quantifiers function exactly the same in both languages. The difference is that one language admits a greater range of logical possibilities than the others. It allows for models that the second language does not. It's this difference which gives rise to the difference in validity, not any difference in the meanings of the premises.

Second, consider the invalidity of identity introduction in free languages. The semantics of identity is identical in both of these languages. $a = b$ is true iff the referent of $a$ is the referent of $b$. Using standard semantics of definite descriptions, if there is an object that is a referent of $a$ and a referent of $b$ and there are no other objects that are referents of $a$ or $b$. In classical languages terms must always refer, and to exactly one object. So $a = a$ is always valid because there is always exactly one referent of $a$, and $a$ obviously co-refers with itself. But in free languages, there are empty terms. Terms are allowed to have to referent. On the assignment where $a$ is empty, $a = a$ is false. There is no object that is the unique referent of $a$ and $a$. Thus, there's a difference in the validity of the rule.

Nevertheless, this difference is not a difference in the semantics of identity. It's a difference in the admittance of models where terms do not refer. The difference occurs at the level of the language as a whole, not at the level of the semantics of particular logical objects. Thus there is a difference in the correctness of the rule without a difference in the meaning of the premises.

More generally, even the cases that prima facie seem to work in favour of the meaning objection can be re-examined. Now, if one compares ordinary classical disjunction to the disjunction of *LP* it's clear that there are differences in truth conditions. They have different truth tables. But consider a language $CL^*$ which contains the connectives of *LP* but where predication is defined in the classical way, with disjoint and exhaustive extensions and anti-extensions of predicates. Classical logic will be the correct logic for this language, even though, strictly speaking, its connectives are those of the language of *LP*. The connectives contain some redundancy, but this is still a perfectly intelligible language. Here, it's far less clear that $\phi \vee_{CL^*} \phi$ has different truth conditions than $\phi \vee_{LP} \phi$. It's simply that the language $CL^*$ rules out some models which are allowed in the language of *LP*. The difference in the validity of disjunctive syllogism is a result of the differences in the models allowed by the two languages, not a difference in the truth conditions of the premises of the argument.

Overall, then, the meaning objection is unsuccessful against both Carnapian views. It misunderstands the kind of pluralism to which the Carnapian is committed. Moreover, it misses a number of possible reasons why there might be differences in valid inference between languages.

## 3.4 Objection 4: The Meta-Logic Objections

### 3.4.1 Quine's Conventions Argument

One of the most historically consequential replies to Carnap's view of logic was Quine (1936)'s argument in *Truth by Conventions*, which I call the Conventions Argument. For decades the argument met with almost universal acclaim. Benacerraf writes:

*"Quine, in his classic paper on the subject, has dealt clearly, convincingly, and decisively with the view that the truths of logic are to be accounted for as the products of [linguistic] convention."* (Benacerraf, 1973, p676)

And if one finds one's self-drawn to the Carnapian position one should:

*"pop a couple of aspirins, re-read your Quine (1936). . . and report back."* (Sider, 2011, p216)

Despite 70 years of broad acclaim, I now take this argument to be sufficiently refuted within recent literature. In particular, Warren (2017) provides a reply that I take to be entirely sufficient[9]. The reply to the argument presented here mirrors many (though not all) aspects of his reply.

Quine's argument goes as follows.

Quine claims that Carnap holds that logical rules are a product of linguistic conventions.

[9]See a number of related works by Warren: *The Possibility of Truth by Convention* (Warren, 2015a), *Talking with Tonkers* (Warren, 2015b), *Change of Logic, Change of Meaning* (Warren, 2018), *Logical Conventionalism* (Warren, 2023) and his book *Shadows of Syntax: Revitalizing Logical and Mathematical Conventionalism* (Warren, 2020)

This is a slightly different presentation from this paper's, but if one views linguistic conventions as the mechanism by which the semantic properties of a language are determined (a highly plausible claim) then my presentation of Carnap and Quine's come together. Conventions determine the semantic properties of a language which in turn determine the nature of truth in the language, which in turn determine the correct logic.

For instance, a classical language might have the convention to only admit terms as grammatical that have a referent, e.g. via a successful baptism. This convention leads to the semantic property that all terms have a single referent, which in turn yields the truth of $t = t$ for all terms $t$ and the validity of identity introduction.

But conventions require a logical framework in which to operate. Conventions have consequences only in relation to a logic. The example of conventions around terms in classical languages, for instance, requires universal instantiation and modus ponens.

A more modern way to put this is that conventions are given in a meta-language but this meta-language will have its own logic.

But then what's to account for the logical rules of the meta-language?
If it's linguistic conventions again, there is a regress.
If it's not linguistic conventions, then the Carnapian view is false.

Warren's reply to Quine is to distinguish between conventions which are explicit and implicit. A convention is explicit if it's overtly introduced during the definition of a language. When building new formal languages, explicit conventions are always used.

The definition of the object language's semantic rules in the meta-language is *prescriptive*. It defines the conventions which determine the semantic properties of the object language. The validity of the logical rules in the object language therefore depend upon the prescription of the conventions in the meta-language. But the conventions depend upon the logic of the meta-language.

Diagrammatically, this is what Quine thinks is going on. Arrows indicating dependence.

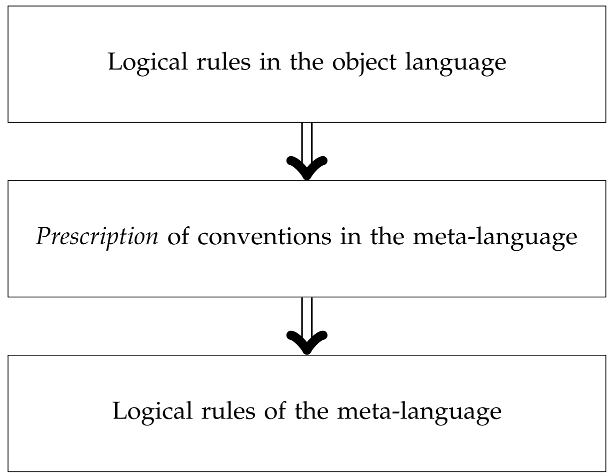


Given this picture, it's clear to see how a regress forms if the logical rules of the meta-language are to be given by a prescription in a meta-meta-language. But this argument assumes a particular relationship between the object language and the meta-language. It's exactly this relationship that Warren's implicit conventionalist targets.

Whilst the conventions and rules of a logic *can* be given explicitly and prescriptively in a meta-language, and for all artificial languages are, they do not need to be given in this way. In natural languages, the semantic properties of the language simply exist as a result of the implicit behaviour and dispositions of the language users. Whilst one can use a meta-language to describe how the speakers of a language behave, the meta-language description of the conventions is merely descriptive, not prescriptive. It states what the conventions and rules of the object language are, but the object language is prior to its meta-linguistic description. The conventions themselves do not depend upon the meta-language in any way.

The implicit conventionalist sees the relationship between the meta-language and object language like this:

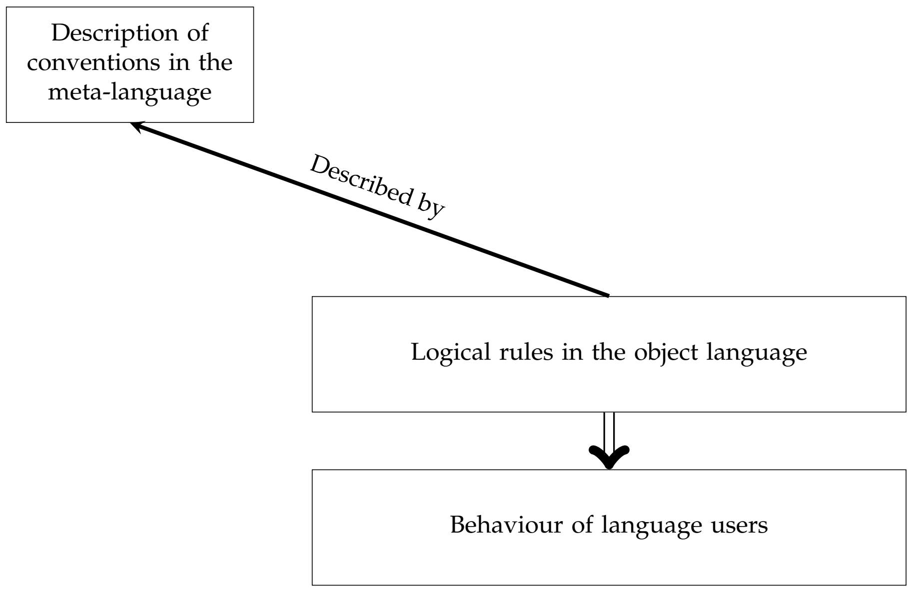


But on this model, there is no regress. One can, of course, inspect the logic of one's meta-theory in a meta-meta-theory, or the logic of that in a meta$^3$-theory, and so on. But the object language doesn't depend on the meta-language on this view.

As an aside, there might also be multiple possible meta-languages with different logics in which we can successfully describe the linguistic conventions and logical rules of some natural language. If there is, this is just Shapiro's logical pluralism.

A natural reply by the Quinian, at this point, is to argue that ordinary linguistic behaviour must be governed by some logic and that in order to learn a language and follow that behaviour one must first understand certain logical notions. This is the adoption problem as discussed by Carroll (1895), Kripke(Kripke (1974a) and Kripke (1974b), Padró (2015), Finn (2019), Cohnitz and Nicolai (2023) and Hattiangadi (2023), amongst others. It's also closely related to Chomsky (1986)'s paradox regarding the acquisition of grammar and Wittgenstein (1953)'s rule-following paradox. But note that the adoption problem is a puzzle regarding how we come to learn logic, not about what the grounds of logical facts are. It's a tangential issue and, though interesting, does not need to be dealt with here.

An important note here is that the Neo-Carnapian does not think that *every* language relates to its meta-theory(s) in the way that Warren's implicit conventionalist describes. Natural languages certainly do. But most artificial languages relate to their meta-theories in the way that Quine describes. The meta-language is used prescriptively; to define the correct rules for the object language. All that is needed to stop Quine's regress, though, is that at some point one can fall back on a natural language whose logical rules are governed by implicit conventions, and hence not subject to Quine's regress.

As a final comment and in continuation from the last point, I wish to clarify a small difference between my own position (the Neo-Carnapian's) and Warren's. Warren uses his solution of Quine's Conventionalism argument as grounds for his general belief in logical inferentialism, the view that the meanings of logical connectives are determined by the inference rules governing them.

The Neo-Carnapian must both express sympathy and disagreement with this view. They acknowledge, pace Warren, that *some* languages have their logical rules determined by implicit conventions. In those languages, inferentialism is likely true. It's the linguistic community's inferential behaviour which determines the correct analysis or analyses of the logical rules of the language. In constructed languages, though, one typically (though not necessarily) starts by defining the semantic rules and deriving the valid logical rules from there. Clearly, in this case, inferentialism is false. The meanings of the connectives are given and the logical rules follow.

### 3.4.2 Logical mismatch between the object and meta-logics

A class of objections near to Quine's are those expressing some kind of displeasure about cases where some language satisfying one logic is either described or defined using a language which satisfies a different logic. This typically occurs when non-classical logics are introduced using a classical meta-theory (Priest, 2008), but in principle, the reverse could happen. This very paper is an example of this phenomenon! I call this logical mismatch, or simply mismatch.

There are many variants of this kind of argument. Space constraints limit the discussion to only a handful. I discuss three that I find most interesting and where I think the replies are most illuminating for the kinds of things that the Neo-Carnapian would say in reply to other variations on this argument.

First, I consider the argument that logical mismatch is simply downright hypocritical. One cannot stand by one set of logical rules one minute and switch to another when it suits.

Second, I consider Griffiths and Paseau (2022)'s variant of this argument, that logical pluralists are committed to an exceptionally weak meta-language.

Lastly, I consider the argument that the meta-language should be, in some sense, privileged with respect to the object language. The logic of the meta-language is, perhaps, more foundational.

The first argument, as discussed by Bacon (2013), is that it's infelicitous, or at the very least embarrassing, to have mismatch[10]. This motivated Bacon's well-known paper on the use of non-classical meta-theory in describing non-classical logics. Strictly, what Bacon shows from a classical meta-meta-theory is that there exists a non-classical meta-theory for non-classical logics. This means that there is still mismatch in Bacon's work. However, with developments in non-classical mathematics over the past decade[11], I take it that there's little doubt to be had that one *could* use a non-classical language to give the meta-theory for its own logic, though I know of no work explicitly doing this.

The objection then is not at non-classical logics per se, as is presented in Bacon's work, but at the logical pluralist who might have a different logic for their meta and object languages. I discuss this from the Neo-Carnapian perspective, but much of what I say also functions as a reply on behalf of logical pluralists more generally.

[10]As an interpretive note, I don't think this version of the argument is quite as prevalent as Bacon thinks. The two main examples he gives, Williamson (1994) and Field (2000), only explicitly oppose mismatch in the very special and specific case of vagueness. I have as of yet been unable to find someone actually advocating for this version of the argument.

[11]This work has mostly been due to Zach Weber, frequently working with other authors: Weber (2010a), Weber (2010b), Weber (2012), McKubre-Jordens and Weber (2012), Meadows and Weber (2016) and Badia et al. (2022). Also see Weber's book *Paradoxes and Inconsistent Mathematics* (Weber, 2021).

The objection, though, only makes sense if one is a logical monist. If one believes that endorsing a logic means endorsing it in all contexts and, crucially, for all languages (or perhaps simply refusing to adopt languages which don't satisfy the one true logic), then it makes sense to be concerned at the presence of mismatch. But without this assumption, it's unclear how the aforementioned discomfort could be motivated. The Carnapian is a logical pluralist. Consequently, they are not worried about using different logics for different circumstances.

A more specific nearby argument is presented by Griffiths and Paseau (2022). This is the argument that Griffiths and Paseau give as a general reply to all pluralist positions, not simply Carnapian pluralism. They contend that if one is a logical pluralist, then one is committed to an exceptionally weak meta-theory. So their argument goes, the meta-theory must be the intersection of all accepted logics. An inference is acceptable in the meta-language iff it is accepted in all correct logics. But as the logical pluralist accepts many logics, their meta-theory will be exceptionally weak.

But Griffiths and Paseau have simply misunderstood what the Carnapian takes the relationship between the object languages and the meta-theory to be. They seem to view the meta-language as a kind of neutral common ground. A domain where users of all logics can come together and unanimously agree on the conclusions. This simply isn't the view of the meta-theory that the Carnapian, or any logical pluralist for that matter, is committed to.

The Neo-Carnapian uses many different logics for many different tasks. They might use a classical logic in mathematics, a constructive logic in computer science, a fuzzy free logic in day-to-day life and quantum logic when discussing fundamental physics. Logics are simply tools for the Carnapian, tools to be picked as appropriate for the task at hand. One such task is studying other logics. The logic best suited for that task might be different from the logic suits to others.

The Neo-Carnapian, then, rejects Griffiths and Paseau's claim that there meta-theory must be the intersection of all correct logics.

This leads directly to the final version of the mismatch objection. Even if the meta-theory does not need to use the intersection of all accepted logics, at the very least the logic of the meta-theory is in some way privileged. The meta-language acts like a judge, passing judgement on what logics can do and what they can't. There's presumably *something* special about that role.

The Neo-Carnapian gives a perhaps surprisingly sympathetic response to this. They agree that meta-theoretic work is exceptionally important. Whilst the job of a meta-language is usually thought of in terms of the kinds of meta-theorems it can prove about some object logic and language (soundness, completeness, etc) another task to be performed in a meta-language is that of determining the appropriateness of a logic for a particular task. Given the Neo-Carnapian's instrumental attitude towards logics, the theory in which they decide which logic (which tool) they use for a particular task is exceptionally important. A craftsperson can have the best tools in the world but if they don't know which to use and when, that's no use. The Neo-Carnapian is then more than happy to accept that the meta-language is exceptionally important, perhaps the most important, language.

This does not, however, undermine their pluralism (either their descriptive pluralism or a corresponding normative pluralism). It's still true that (1) there are logics other than the one correct for their meta-language that are satisfied by other languages and (2) there are logics other than the one correct for their meta-language which are appropriate for use in some context or to some end. Thus the privileged place that the Neo-Carnapian gives their meta-theory is not one that undermines their pluralism.

As a final comment, I talk of *the* meta-language when, in fact, there are many. There are many meta-languages which could perform the task of adjudicating which logic should be used in which situation. One of two things is the case: (1) even if there are many languages which *could* function as the right language for logic selection, there is one that is the best. The Neo-Carnapian would use that one. (2) There are many languages which are equally best suited to the task of logic selection, in which case the Neo-Carnapian is happy with anyone using any of these, even if they yield differing results[12].

[12]There is also a possible third option if there are many orthogonal dimensions of evaluation for theories of logic selection. There might be some languages/logics best under some dimensions of evaluation but others best qua others. If that is the case, I don't know what the Neo-Carnapian should say, though I'd take that to be an exceptionally interesting case if it were to arise.

# 4 Conclusion

This paper defends the much-maligned Carnapian account of logical correctness and Logical Pluralism. In particular, I outline two possible Carnapian accounts (§1): The classical Carnapian account and my own Neo-Carnapian view. After clarifying the type of Logical Pluralism that the Carnapian is committed to (§2), four well-known and influential objections to the position are outlined and my Neo-Carnapian's replies explained (§3). In particular, I reply to Cook (2010)'s objection that Carnapian logical pluralism is insubstantial (§3.1), Prior (1960)'s Tonk objection (§3.2), Quine (1970) and Restall (2002)'s meaning objection (§3.3) and various meta-logical objections, including Quine (1936)'s famous conventions objection (§3.4.1) and the mismatch objections from, amongst others, Griffiths and Paseau (2022) (§3.4.2). I've shown, I hope, that a version of the Carnapian view, my Neo-Carnapian view, survives the standard objections levied against it.

# References

Bacon, A. (2013). Non-classical metatheory for non-classical logics. *Journal of Philosophical Logic*, 42(2):335–355.

Badia, G., Weber, Z., and Girard, P. (2022). Paraconsistent metatheory: New proofs with old tools. *Journal of Philosophical Logic*, 51(4):825–856.

Beall, J. and Restall, G. (2005). *Logical Pluralism*. Oxford, England: Oxford University Press.

Belnap, N. D. (1962). Tonk, plonk and plink. *Analysis*, 23:130–134.

Benacerraf, P. (1973). Mathematical truth. *Journal of Philosophy*, 70(19):661–679.

Blake-Turner, C. and Russell, G. (2018). Logical pluralism without the normativity. *Synthese*, pages 1–19.

Carnap, R. (1928). *Der Logische Aufbau der Welt*. Hamburg: Meiner Verlag.

Carnap, R. (1934). *Logische Syntax der Sprache*. Springer, Vienna. Translated into English 1937 by Amethe Smeaton under the title "Logical Syntax of Language", Routledge, London.

Carnap, R. (1947). *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago, IL, USA: University of Chicago Press.

Carnap, R. (1950). Empiricism, semantics and ontology. *Revue Internationale de Philosophie*, 4(11):20–40.

Carroll, L. (1895). What the tortoise said to achilles. *Mind*, 4(14):278–280.

Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origin, and Use*. Prager.

Cohnitz, D. and Nicolai, C. (2023). How to adopt a logic. *Dialectica*.

Cook, R. T. (2010). Let a thousand flowers bloom: A tour of logical pluralism. *Philosophy Compass*, 5(6):492–504.

Field, H. (2000). Indeterminacy, degree of belief, and excluded middle. *Nous*, 34(1):1–30.

Finn, S. (2019). The adoption problem and anti-exceptionalism about logic. *Australasian Journal of Logic*, 16(7):231.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23(6):121–123.

Griffiths, O. and Paseau, A. C. (2022). *One True Logic: A Monist Manifesto*. Oxford University Press.

Hattiangadi, A. (2023). Logical conventionalism and the adoption problem. *Aristotelian Society Supplementary Volume*, 97(1):47–81.

Kouri, T. (2019). A new interpretation of carnap's logical pluralism. *Topoi*, 38(2):305–314.

Kripke, S. (1974a). Princeton lectures on the nature of logic. Transcription.

Kripke, S. (1974b). The question of logic. Transcription of a lecture given at the University of Pittsburgh.

McKubre-Jordens, M. and Weber, Z. (2012). Real analysis in paraconsistent logic. *Journal of Philosophical Logic*, 41(5):901–922.

McSweeney, M. M. (2019). Following logical realism where it leads. *Philosophical Studies*, 176(1):117–139.

Meadows, T. and Weber, Z. (2016). Computation in non-classical foundations? *Philosophers' Imprint*, 16.

Padró, R. (2015). *What the Tortoise Said to Kripke: The Adoption Problem and the Epistemology of Logic*. PhD thesis, CUNY.

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

Prior, A. N. (1960). The runabout inference-ticket. *Analysis*, 21:38–39.

Quine, W. V. (1936). Truth by convention. In *Philosophical Essays for Alfred North Whitehead*, pages 90–124. London: Longmans, Green & Co.

Quine, W. V. O. (1970). *Philosophy of Logic*. Cambridge, Mass.: Harvard University Press.

Restall, G. (2002). Carnap's tolerance, meaning, and logical pluralism. *Journal of Philosophy*, 99(8):426–443.

Russell, G. (2020). Logic isn't normative. *Inquiry: An Interdisciplinary Journal of Philosophy*, 63(3-4):371–388.

Shapiro, S. (2014). *Varieties of Logic*. Oxford and New York: Oxford University Press.

Sider, T. (2011). *Writing the Book of the World*. Oxford University Press, Oxford.

Soames, S. (2003). *Philosophical Analysis in the Twentieth Century: Volume 1: The Dawn of Analysis*. Princeton University Press, Princeton.

Tahko, T. E. (2021). A survey of logical realism. *Synthese*, 198(5):4775–4790.

Warren, J. (2015a). The possibility of truth by convention. *Philosophical Quarterly*, 65(258):84–93.

Warren, J. (2015b). Talking with tonkers. *Philosophers' Imprint*, 15.

Warren, J. (2017). Revisiting quine on truth by convention. *Journal of Philosophical Logic*, 46(2):119–139.

Warren, J. (2018). Change of logic, change of meaning. *Philosophy and Phenomenological Research*, 96(2):421–442.

Warren, J. (2020). *Shadows of Syntax: Revitalizing Logical and Mathematical Conventionalism*. New York, USA: Oxford University Press.

Warren, J. (2023). Logical conventionalism. In Brendel, E., Carrara, M., Hjortland, O., Sagi, G., Sher, G., Steinberger, F., and Ferrari, F., editors, *Oxford Handbook of Philosophy of Logic*. Oxford: Oxford University Press.

Weber, Z. (2010a). Extensionality and restriction in naive set theory. *Studia Logica*, 94(1):87–104.

Weber, Z. (2010b). Transfinite numbers in paraconsistent set theory. *Review of Symbolic Logic*, 3(1):71–92.

Weber, Z. (2012). Transfinite cardinals in paraconsistent set theory. *Review of Symbolic Logic*, 5(2):269–293.

Weber, Z. (2021). *Paradoxes and Inconsistent Mathematics*. Cambridge University Press.

Williamson, T. (1994). *Vagueness*. Routledge.

Wittgenstein, L. (1953). *Philosophical Investigations*. New York, NY, USA: Wiley-Blackwell.

# Essay VI: Three Approaches to Logical Correctness

## Introduction

This paper outlines three different ways one might think about logical correctness: the Realist approach, the One-Langauge approach and my own Neo-Carnapian view. The realist and one-language views have dominated the philosophy of logic in recent years. I argue against these approaches and in favour of the Neo-Carnapian approach.

The Realist view holds that logical correctness is language and mind-independent.

They hold the following:

**OBJ:** There are logical facts and they are objective, in the sense that they are mind and language-independent.

**REL:** These facts matter for logical correctness. The interesting notion of logical correctness is correctness with respect to the logical facts.

I have in mind individuals like Frege (Frege (1879) andFrege (1893)), the early Wittgenstein (1922); Tahko (Tahko (2014) and Tahko (2021)) and McSweeny (McSweeney (2018) and McSweeney (2019)), from whom I take the term term 'logical realist', along with naturalists like Quine (Quine (1960), Quine (1970) and Quine (1981)) and Maddy (Maddy (2007) and Maddy (2012)).

The One-Language view takes logical correctness to be semantic and hence not language-independent. However, what they take to be interesting for the philosophy of logic is correctness for some fixed language.

**SEM:** Logical truths are semantic in that they are mind but not language-independent.

**OL:** The semantic truths, given some fixed language, are what matter for logical correctness. The interesting sense of logical correctness is correctness with respect to the semantic facts of some fixed language.

I have in mind individuals like Dummett (1991), Wright (1992), Pedersen (2014), Beall and Restall (2005), Shapiro (2014), Cook (Cook (2010) and Cook (2023)) and Griffiths and Paseau (2022). Cook, in particular, has been quite vocal in calling versions of logical pluralism which don't fix a language "insubstantial".

Within One-Language views, I also include individuals like Lewis (Lewis (1986) and Lewis (1998)) who engage in linguistic explication. This is discussed below. Others on the above list could also be categorized along with Lewis as explicators.

Lastly, the Neo-Carnapian. This view is my own and, whilst inspired by Carnap (1934)'s view and sharing many similarities, breaks from it at a crucial juncture. The Neo-Carnapian approach agrees with **SEM** but rejects **OL**. They are interested in how logics can be correct across a range of languages, not simply their own. Moreover, the Neo-Carnapian takes there to be a complex web of reasons speaking for or against the adoption of one language or another. They see this web of reasons as of great importance to the philosophy of logic. The Neo-Carnapian has two notions of logical correctness. First, a logic can be descriptively correct for some language iff it captures the truth-preserving inferences of that language. Second, logics can be normatively correct if they are descriptively correct for a language one has reason to adopt.

The Neo-Carnapian approach claims the following:

**SEM:** Logical truths are semantic in that they are mind but not language-independent.

**ML:** The semantic truths across many languages are what matter for descriptive logical correctness. A logic is descriptively correct for some language iff it captures the valid inferences in that language.

**NORM:** There exists a complex web of reasons for and against adopting one language over another. A language is normatively correct iff it is descriptively correct for a language those reasons speak in favour of.

Here I have in mind my own view, though this is greatly inspired by Carnap's view.

§1 clarifies some preliminaries. §2 presents Carnap's view, then contrasts it with my own Neo-Carnapian view. §3 outlines the realist approach and gives two reasons against it. A third reason is mentioned in passing but not discussed in detail. §4 outlines the one-language approach, gives two reasons against it and defends the Neo-Carnapian view from an objection.

# 1 Preliminaries

## 1.1 Logic as the normative laws of reasoning?

There's a fourth possible approach to logical correctness that can be found in the work of, amongst others, Frege and Priest (Priest (1987), Priest (2006) and Priest (2010), amongst other places): the normative approach.

On this view logic is the study of the laws of reasoning. Now, what's meant here is not the *descriptive* laws of reasoning, that would be Psychologism and is, I think, close to universally rejected. Logic is the study of the *normative* laws of thought, i.e. the study of how we *should* reason, the kinds of inferences we may or ought make. On this view, a logic is correct iff it correctly describes the normative laws of reasoning.

Some version of this view is obviously true. Clearly, something logics needs to do is adjudicate the success of arguments. Philosophy of logic must include the philosophy of that adjudication.

However, accepting this claim doesn't speak for or against any of the three views discussed in this paper. Language and reasoning are clearly connected. Language is typically the vehicle for reasoning or, at least, the way in which reasons are *expressed*. It makes sense, then, for language to be the direct object of critique of logic, as linguistic inferences can actually be evaluated in a way pure thoughts cannot, even if this is just a proxy for "proper" reasoning.

Blake-Turner and Russell (2018), and Russell (2020) as a solo author, have replied to this view, arguing that logic should be thought of as the descriptive study of truth-preserving inference rather than the normative study of reasoning. Given that reasoning is concerned with truth preservation, a descriptively true theory of truth preservation is also a normatively true theory of reasoning.

I'm convinced by their explanation for two reasons.

First, most logic textbooks don't discuss normativity but do discuss truth. For example, Halbach (2010)'s "Logic Manual" contains the word "truth" 254 times but fails to use the word "norm" or any co-conjugations thereof even once. Even Priest, an advocate of the normative view of logic, never uses the word "normative" in his non-classical logics textbook (Priest, 2008). The word "norm" appears once, though not in connection to reasoning. "Truth" appears 64 times. This doesn't speak *decisively* in favour of Russell's claim. It might just be that truth talk is more pedagogically useful than norm talk. But this at least provides some prima-facie evidence that logic is really about truth-preserving inference and merely used as a normative tool.

Second, there are many norms of reasoning that aren't of interest to logic. One might have obligations to one's friends or loved ones to believe certain things about them. One might have duties to one's self to believe in a way that promotes what's best for you (James (1896)). One might have obligations not to believe racist beliefs (See the collection by Kim and McGrath (2019), in particular, Basu and Schroeder (2019)). All of these are norms of reasoning but none are of interest to logic. Russell's account explains why only some norms of reasoning are relevant to logic.

Given, then, that the normative view of logic collapses into Russell's descriptive account, what matters for logic is truth preservation. Logical monism/pluralism/nihilism is then true iff there are is one/many/no correct theory/theories of truth preservation.

But one might be a realist about truth, claiming that there are language- and mind-independent facts about the nature of truth (Wittgenstein (1922),Armstrong (1997) and Maddy (2007), amongst others). Alternatively, one might think truth is semantic (Tarski (1931), Tarski (1943) and Field (2001)). But if truth is semantic, should one be concerned with truth in some fixed language or how truth differs across languages?

Whichever answers one agrees with, one arrives at one of the three positions outlined above.

## 1.2 What do the three approaches disagree on?

I don't believe that there's an objectively correct notion of logical correctness. The disagreement between the three approaches outlined is not about which notion of logical correctness, and hence which versions of logical monism, pluralism and nihilism, is correct, in some objective sense. All three approaches make perfectly intelligible claims about the relationship between logics, languages and reality. It's uninteresting if the ordinary notion of "logical correctness" is closer to Realism, the One-Language view or the Neo-Carnapian view.

The disagreement about the relative importance of these three competing notions. The three approaches might grant that each other's notions of correctness are perfectly intelligible, whilst rejecting that this is interesting or important for logic or the philosophy of logic.

Suppose for instance that the following is the case:

1. A logic $L_1$ perfectly captures the logical structure of reality.
2. A logic $L_2$ perfectly captures the semantics of our actual language.
3. A logic $L_3$ is the most advantageous logic to work with, in a given context.

The logical realist, one-language theorist and Neo-Carnapian all know these facts.

The logical realist takes (1) to speak decisively for $L_1$. After all, $L_1$ is objectively right. The world is the way $L_1$ says it is.

The one-language theorist takes (2) to speak decisively for $L_2$. It doesn't matter that $L_1$ is 'right' in some objective sense. We don't reason in the pure concepts of the universe, we reason in our actual potentially flawed language. Consequently what matters is if our actual language endorses an inference.

The Neo-Carnapian takes (3) to speak decisively for $L_3$, within this context. (1) might be interesting; there are certainly some contexts where capturing the logical structure of reality is important. (2) is just irrelevant. We can change our language quite easily; there's no good reason for conservatism for conservatism's sake. If $L_2$ is sub-optimal, our language should just be changed to satisfy a better logic. What matters is understanding what a context requires our logic to do (which might include capturing the logical structure of reality) and picking a logic that best meets those requirements. $L_3$ is that logic.

What this means is that the disagreement between the three approaches isn't per se about which logics relate to which kinds of linguistic or metaphysical structures in which ways, but rather about what those relations mean for the adoption of a particular logic. The disagreement is about the philosophical upshot of potentially agreed-upon facts, not necessarily the facts themselves.

Now, they might also disagree on these facts. In particular, there are plausible objections to Logical Realism that deny that there is even such a thing as the objective logical structure of reality, or that we could know about it. §3.2 briefly discusses these objections. The point is that one can just as easily object to logical realism on the grounds that objective logical facts are irrelevant as one can that there are no objective logical facts[1].

[1]Even if one were to show that there are no objective logical facts, that wouldn't necessarily mean that logical realism is wrong. It might just mean that logical nihilism is true - logics are correct iff they are descriptively true theories of the objective logical facts. There are no objective logical facts. So there are no correct logics.

# 2 The Carnapian & Neo-Carnapian approaches to Logical Correctness

## 2.1 Carnap's Philosophy of Logic

Carnap agrees with **SEM** and **ML**. He takes logics to be correct relative to a language. Logical facts are mind but not language-independent. The majority of LSL is devoted to showing how a first-order logic and a type theory could be true of two different languages.
For example, one might define a language where predication is Boolean or a language where it allows for predication gluts or gaps. This would then yield a difference in the validity of certain logical rules. Disjunctive syllogism, for instance, is valid in two-valued logics but not glutty three-valued logics[2]. One might define a language where all terms refer or where they may be empty terms. These semantic choices lead to different correct logics in that language.

However Carnap, at least in some places, appears to disagree with **NORM**. His tolerance principle states the following:

> *"In logic there are no morals. Everyone is at liberty to build up [their] own logic, i.e. [their] own language. All that is required of [them] is that, if [they] wishes to discuss it, [they] must state [their] methods clearly, and give syntactical rules instead of philosophical arguments."* LSL P52

He continues saying that: *"It is not [the business of philosophy of logic] to set up prohibitions, but to arrive at conventions"*.

On Carnap's view, that a logic is correct for some language is sufficient to reason with it, provided one is clear about one's choice of language. There's little to no grounds to choose between competing logics, just a matter of convention. There's certainly not the Neo-Carnapian's web of reasons as outlined in **NORM**.

[2]Consider the logic *LP*. There are three truth values: $T$, $F$ and $B$. $T$ and $B$ are designated. $\phi \vee \psi$ takes the maximum value of $\phi$ and $\phi$ on the order $T > B > F$. $\neg\phi$ behaves classicalls for $T$ and $F$ and takes $B$ to itself. Now consider the assignment where $\phi$ is $F$ and $\psi$ is $B$. $\phi \vee \psi$ and $\neg\psi$ are $B$, both designated, but $\phi$ is $F$, undesignated. Thus disjunctive syllogism is invalid in this logic.

There are, I believe, some parts of Carnap's wider work that perhaps suggest a more nuanced view here. ESO is all about choosing between nominalist and non-nominalist languages, so Carnap clearly does think there are some grounds on which one might choose between rival languages[3]. Moreover, one might hope that the process of arriving at conventions is not arbitrary; it should have some structure or rules. On this reading of Carnap, the gap between Carnap and the Neo-Carnapian might be smaller than I present it here, or even non-existent. I leave it to more able Carnap scholars than myself to clarify this point.

For present purposes, I'm going to take Carnap's view to be his view in LSL, which accepts **SEM** and **ML** but rejects **NORM**. He does not take there to be reasons in favour of or against adopting one logic over another. The Neo-Carnapian sees an interesting and complex network of reasons pertaining to language, and hence logic, selection.

To put this more simply: Carnap accepts his tolerance principle so cannot accept **NORM**. The Neo-Carnapian agrees with Carnap about the nature of logical truth (**SEM** and **ML**) but rejects the tolerance principle in favour of **NORM**..

## 2.2 The Neo-Carnapian view explained

Like Carnap, the Neo-Carnapian takes logical correctness to be a language-relative matter. A logic is correct only if it's sound. Typically, a logic will be correct iff it is sound and complete, but this requirement will obviously have to be weakened in the context of higher-order logics.

However, soundness and completeness depend on the nature of truth and the nature of truth varies from language to language. Languages might differ in the connectives they use, the number of truth values they have, the nature of identity, predication or quantification, or any other number of determinables. Consequently, which logic is correct varies from language to language.

[3] In the concluding remarks of ESO Carnap says the following: *"The acceptance or rejection of abstract linguistic forms, just as the acceptance or rejection of any other linguistic forms in any branch of science, will finally be decided by their efficiency as instruments, the ratio of the results achieved to the amount and complexity of the efforts required."*

The relevant question, then, is how to decide which language one ought work in. For Carnap, at least as I'm presenting him, there's not much to say here. It's simply a matter of arriving at a convention, there are no stronger normative considerations which determine which language one really *should* work within.

For the Neo-Carnapian, however, this is where one of the most interesting parts of the philosophy of logic begins. The Neo-Carnapian sees a complex web of reasons speaking for or against one language or another. There are questions about how one might want predication or naming to function, or what truth, falsity or other truth values should signify about the true statement. There are questions about what we want our logics to do or how we want our language to go about carving up the world.

One of the central tasks of the philosophy of logic, according to the Neo-Carnapian, is to map and understand these reasons and how they relate to one another. Language selection, for the Neo-Carnapian, is then a complex and philosophically rich issue.

To return momentarily to §1.1, I think it's helpful to consider how the Neo-Carnapian view thinks about the norms of reasoning to understand their view. The Neo-Carnapian has a two-tiered view of the norms of reasoning. On the higher tier, there are norms relating to language selection. There are reasons that speak for or against the use of one language or another. There are then lower-tier reasons that say, given some choice of language, how one ought reason. These lower-tier reasons are, to use Kant's terminology, hypothetical imperatives.

In summary, then, the Neo-Carnapian makes the following claims:

**SEM:** Logical truths are semantic in that they are mind but not language-independent.

**ML:** The semantic truths across many languages are what matter for descriptive logical correctness. A logic is descriptively correct for some language iff it captures the valid inferences in that language.

**NORM:** There exists a complex web of reasons for and against adopting one language over another. A language is normatively correct iff it is descriptively correct for a language those reasons speak in favour of.

### 2.3 A brief comment on reasons pertaining to logic selection

Having stated that there the primary difference between Carnap's view and the Neo-Carnapian view is the acceptance of a network of normative reasons counting or against one choice of logic or another, it's important to provide a little more detail as to what these reasons look like.

I do not wish to give an exact theory of the reasons pertaining to logic selection here. Different neo-Carnapian views might differ on exactly what the correct normative theory of logic selection is. I personally haven't settled on a singular account. It would therefore be unhelpful to commit the Neo-Carnapian approach to a singular normative theory at this stage. That being said, it's worth considering the kinds of considerations that might appear in these sorts of theories.

Context must play a large role in any good normative theory of logic selection. Contexts, amongst other things, often contain certain goals or aims. The context of fundamental physics, for instance, might have the goal of capturing the structure of reality as closely as possible. Consequently, the kinds of metaphysical consideration outlined below in §3 might be more important within this context. Many social contexts might need to track many types of identity across time and might therefore benefit from the inclusion of non-rigid designators. Many contexts in the macroscopic world do not need to specify every vague boundary and hence might benefit from adopting a fuzzy logic.

It might even be the case that all of the norms of logic selection come from contextual factors. There might be no "universal" norms of logic selection. I remain neutral on this question. There's certainly something to be said for it, but I don't think a decisive case can be made in favour of it.

Other considerations might be pragmatic. Some logics might simply be easier or more productive to reason with. For instance, all else being equal, a stronger logic is likely better. If a logical rule could be unproblematically adopted within a context without undermining some other beneficial feature, then it likely should be.

Whilst this is far from a full normative theory of logic selection, this hopefully clarifies the kinds of considerations that the Neo-Carnapian takes to be relevant to logic selection.

# 3 The Realist approach to Logical Correctness

## 3.1 Logical Realism Explained

Logical Realism holds that there are objective logical facts. Logic is the descriptive study of those facts. It's modern defenders include the likes of Tahko (Tahko (2014) and Tahko (2021)), McSweeny (McSweeney (2018) and McSweeney (2019)) and Maddy (Maddy (2007) and Maddy (2012)). Historic defenders include but are certainly not limited to Frege (Frege (1879) and Frege (1893)), the early Wittgenstein (1922) and Quine (Quine (1960) and Quine (1981)).

I take Logical Realism to amount to the following two claims:

**OBJ:** There are logical facts and they are objective, in the sense that they are mind and language-independent.

**REL:** These facts matter for logical correctness. The interesting notion of logical correctness is correctness with respect to the logical facts.

What are logical facts and what are they like?

Exactly what one wants to say about logical facts will depend on one's metaphysical views. Tahko, for instance, is a committed Neo-Aristotelian and would likely have something quite different to say than, say, Wittgenstein. It would be more than a little impractical to enumerate all the possible metaphysical positions and how they might go about adopting the metaphysical approach. what I present instead is an example position that serves as an exemplar of this kind of position, a dummy model that captures the essence of the metaphysical approach.

Although the arguments of this section are presented in response to the exemplar, this is only for ease and brevity of explanation. The arguments will generalise quite naturally to other versions of logical realism as well, though I want to leave open the option that especially well-crafted versions of logical realism might find inventive ways of avoiding one, many or all of my objections.

For the exemplar, I use a factive ontology, which has elsewhere been called a truthmaker ontology or an assertory ontology. On this view, the world is, at least in part, composed of facts. By facts, here, I do not simply mean true sentences. No one denies that there are facts in the sense of true sentences. What the factive ontologist asserts is that there are language-independent *things* called facts.

Exactly what those facts are will differ between different versions of this view, but a natural view would be that they're abstract objects. Sentences are about facts, perhaps some sentences about many facts. They are true iff the fact(s) they are about exists[4].

Facts stand in certain relations to one another. A fact $\phi$ might be the negation of a fact $\psi$. A fact $\rho$ might be the conjunction of the facts $\phi$ and $\phi$.

A basic category of facts are what we might call first-order facts. These are simple descriptive facts about the world[5]. These include atomic facts, which predicate a relation of some objects, truth-functional combinations of first-order facts, such as conjunctions, disjunctions, etc, and quantified statements about the world, such as *"all men are mortal"* or *"there are some critics who only admire one another"*.

There are also facts about facts, I'll call these higher-order facts. For instance, there might be the fact that, necessarily, it's never the case that both a fact and its negation exist. This is the metaphysical formulation of the law of non-contradiction -**LNC**.

Logical facts are higher-order facts about what possible facts can exist alongside others. For instance, if it might say that if the fact that $\phi$ exists then the fact $\phi \vee \psi$ exists as well. Logical facts take the form "If the facts $x_1$, $x_2$, $x_3$, etc do/don't exist; then the facts $y_1$, $y_2$, $y_3$, etc do/don't exist".

On the simplest version of logical realism, a logic is correct iff it's a correct descriptive theory of the logical facts[6]. So if the existence of facts $x_1$, $x_2$ and $x_3$ necessitate the existence of fact $y$ then it should be the case that $x_1, x_2, x_3 \vdash y$ in the logic.

[4]There's an alternative version of this view where all facts exist necessarily but contingently have the properties of true, false or whichever other truth values there might be.

[5]This is not to be confused with first-order in the *quantificational* sense. I'm contrasting facts about the world with facts about facts, not different types of quantification.

[6]Or at least as many of them as possible, given completeness constraints in higher order logics

There are more fined-grained versions of logical realism that are even more demanding. McSweeny (McSweeney (2018) and McSweeney (2019)), for instance, argues that there are certain privileged sets of logical operators. Universal quantification might be more basic than existential quantification, for instance. On a factive view, this would correspond to a fundamentality relation between facts. The universal fact that $\forall x\phi$ is more fundamental than the existential fact that $\neg\exists x\neg\phi$.

Resolving the argument between McSweeny and more modest logical realists isn't necessary for this paper and the point at hand. What's worth noting for present purposes is that this debate essentially amounts to an argument about what the logical facts are. The view presented here is just concerned with what facts exist. McSweeny argues that the logical facts also include facts about quite fine-grained dependence between logical facts and that these also need to be reflected in the logic. The outcome of this debate isn't necessary for this paper, though, so no more is said on it in the main paper. A few comments are made in footnote, however[7]

How does this work in practice when considering debates between rival logics? Here are two examples: fuzzy vs classical logic, and free vs classical logic.

[7]I do think McSweeny's version of logical realism is less plausible than more modest views. McSweeny needs to establish two things: (1) There are facts about the objective priority between logical connectives, e.g. which complete sets of connectives are *really* fundamental and which are *really* derived (2) That we should actually care about these facts, if they do exist.

On (1), I confess some level of confusion. I'm honestly unsure what the world would have to be like in order for, say, existential quantification to be objectively more fundamental than universal quantification. What are the ways the world could be that would make either of these true?

There's also an epistemological objection to (1). Even if one of the answers is true, how could we know? But McSweeny is open to these being facts that we could never know, so I won't push further here.

A more pressing objection, though, is (2). Even if there's an objective priority between connectives, for instance, it's not obvious why one should care about that when selecting one's logic. It's unclear how we're better equipped to understand the world by using a logic which respects the objective priority between, say, the logical operators.

Consider the debate over fuzzy logic. On non-fuzzy views, predication is binary. Objects satisfy a predicate or they don't. Metaphysically, this amounts to the claim that there exist full facts about some objects satisfying the predicate (or negative facts about them failing to). In fuzzy logic, objects can partially satisfy predicates. It can be, say, 20% true that $P(a)$. Metaphysically, this could amount to the claims that (1) there exist partial facts - i.e. facts about partial states of affairs or (2) full facts can partially exist.

But these are both claims about the logical facts. The question about which logic is correct reduces to a set of claims about the nature of these facts.

Similarly so in the debate between classical various types of free logic. Free logics break from classical logics in that they allow for terms which don't refer, e.g. when the referent doesn't exist. Free logics can be negative if atomic sentences containing empty terms get false truth values or gappy if atomic sentences containing empty terms get no truth value.

There's some question as to how classical logic should handle non-existence. Terms in classical logic have to refer. Certainly, if we have classical logic with rigid designators, necessitism follows, i.e. everything that possibly exists necessarily exists. Assume for the present case that names do have to be rigid designators, though an analogous case challenging this is considered in §3.3.

If necessitism is true, then classical semantics for terms are correct. For any object $o$, $o$ exists in every world, so the fact that $o = o$ exists in every world, so classical identity introduction is valid.

If necessitism is false, again assuming names have to be rigid designators, then the classical semantics for identity cannot be true. There are contingent objects $o$ which don't exist in some worlds. In those worlds the fact $o = o$ doesn't exist. So classical identity introduction is invalid.

Exactly which free logic would be correct depends on if the fact that $o \neq o$ exists in worlds where $o$ does not. And, more generally, if negations of atomic facts about $o$ exist in non-$o$ worlds. If they do, negative free logic is true. If they don't, gappy-free logic is true.

The crux, though, is that all of these are questions about the logical facts which, according to logical realism, determine logical correctness.

In summary, the logical realist claims the following:

- There are objective (i.e. language and mind-independent) logical facts.
- The logical facts are higher-order facts about the relationship between other kinds of facts
- A logic is correct iff it correctly describes the logical facts. There are some different options available as to which facts are the logical facts.
- Given this, logical questions are resolved by metaphysical questions.

I discuss three objections.

## 3.2 Standard Nominalist Objections

Familiar nominalist objections against this kind of position can be raised. Following Burgess and Rosen (1997), there are three types of familiar nominalist objections: Semantic, Ontic and Epistemic.

The semantic argument, following someone like Ayer (1936), might argue that talk of strange metaphysical objects called facts is, simply put, non-sense.

Similarly one might argue for a world where there are no abstract metaphysical facts, just ordinary concrete objects.

Lastly, one might argue that even if there are things called facts, one could never know about them (Benacerraf (1965) and Benacerraf (1973)).

Each of these three objections provides a different way to reject **OBJ**.

This is, though, all rather well-trodden ground, at least with respect to the analogous issue of mathematical realism.

To take the epistemic objection to the metaphysical view, for instance, one might reply to the objection that logic plays an indispensable role in our best current science (Quine, 1981) or best scientific explanations (Lange, 2016). This would give us knowledge of the logical facts. But perhaps the role of logic in explanation is instrumental, part of an idealization (Leng (2021)). In this case, scientific knowledge would not give us logical knowledge.

Perhaps the epistemic objection is based on a faulty causal theory of knowledge (Liggins, 2006). Perhaps any theory of knowledge should take logical knowledge as a known data point and hence any epistemic objection must rely on a false theory of knowledge (Lewis, 1993).

To add an extra level of complexity here, not all of these arguments will generalize to all versions of logical realism. These arguments are really arguments against abstract objects, so work as well as they always do against the factive ontology I present here, but might entirely miss their mark with something like Tahko's Neo-Aristotelian logical realism.

So for four reasons, I'm not going to discuss this line of objection any further:

1. These are well-trodden debates about which I couldn't hope to briefly say anything novel
2. I don't think these objections generalize to all versions of logical realism, or at least they will vary in detail so heavily that I couldn't approach them uniformly.
3. These arguments involve highly contentious claims about the nature of knowledge, language and reality. I don't want to hang the case against the metaphysical approach on such contentious assumptions.
4. Even if there are no facts (in the metaphysical sense) or we can't know about them, that doesn't rule out logical realism. Logical nihilism or logical scepticism might be true.

That's not to say the standard nominalist arguments aren't good arguments, they might even be sound. They're just not the ones most useful for present purposes.

Instead, I give two hopefully less contentious and perhaps more novel arguments against logical realism. I target **REL** rather than **OBJ**. First, even taking logical realism at its strongest, there are debates in the philosophy of logic which aren't resolved by objective logical facts. This means that logical correctness can't *merely* be a matter resolved by the logical facts. Metaphysics might play a role, but it requires a little help. Second, even if there are logical facts, it's unclear why they are always relevant to logic selection and a notion of logical correctness not relevant to logic selection is less relevant than one that is.

### 3.3 Logical facts don't resolve every debate between logics

This section argues that even taking logical realism at its strongest, there are still questions about the validity of certain logical rules which are not resolved by the logical facts alone. For the sake of argument, grant, contra the previous section, that there are logical facts and that we have perfect epistemic access to these facts. Nevertheless, there are some questions about the validity of logical rules which are not fully resolved.

Consider the logical rule that allows for the substitution of identity inside modal operators. Call this the **MSR** for the modal substitution rule.

The prima facie grounds for accepting or rejecting **MSR** are clear for the logical realist. Assume, as in the case of free vs classical logics above, that names are rigid designators.

**MSR** is true iff if two terms refer to the same thing in some world, they refer to the same thing in all worlds (where either refers at all). Assuming that terms are rigid designators, this is true iff objects that are identical in one world are identical in all worlds (necessity of identity). This is a substantive metaphysical claim that might be true or false, but what matters is that the status of **MSR** is determined by the logical facts about the necessity of identity.

The issue with this argument is that the assumption that names are rigid designators does a lot of work. Names certainly don't *have* to be rigid designators. One can make perfectly good sense of languages containing non-rigid designators.

An example of a name that might quite naturally be understood as a non-rigid designator is a name like 'Spiderman'. To illustrate a point, I ignore the distinction between fictional and non-actual names. I'll assume that everything that takes place within the fictional story of Spiderman takes place in some possible world and consider the semantics of the name 'Spiderman' within that world.

'Spiderman' is not a rigid designator. At some times in some worlds, 'Spiderman' refers to Peter Parker. At other times in other worlds, 'Spiderman' refers to Miles Morales. There are some worlds where it refers to neither. It changes its reference and is therefore a non-rigid designator.

Of course, one could still work in a language where names are rigid designators and accommodate for 'spiderman' in known ways. There are two options: (1) When the apparent reference change happens, one could hold that this is actually the baptism of a new name. There are the two distinct names $S_{mm}$ and $S_{pp}$ (2) One could hold that "Spiderman" is really just a (definite) description in disguise, akin to designators like "The King of France".

I'm certainly not claiming that a name like 'Spiderman' can only function as a non-rigid designator. I'm simply claiming that it would be perfectly possible to have a language where 'spiderman' is a non-rigid designator. I think it's even likely that English is one such language, but that's contentious.

Why does this matter for logical correctness?

**MSR** is invalid in languages with non-rigid designators. Necessarily, Spiderman is Spiderman. Contingently, Spiderman is Miles Morales. It would follow by **MRS** that necessarily Miles Morales is Spiderman, but this would be false.

Note that the presence of non-rigid designators is a sufficient but not necessary condition for **MSR** failing. If necessity of identity, the metaphysical law, is false, **MSR** fails either way. **MSR** only holds when (1) necessity of identity is true and (2) all names are rigid designators.

What does this mean for logical realism?

It's a semantic not a metaphysical matter whether names can be non-rigid designators in some particular language. The logical facts do not determine if a language contains non-rigid designators. But **MSR** depends on the claim that names are rigid designators. So **MSR** depends, at least in part, on something other than the logical facts. This means that the correctness of a logic containing **MSR** is not entirely determined by the logical facts.

Three objections on behalf of the logical realist.

Objection 1: We know from Kripke (1980) that names have to be rigid designators.

Reply 1: Kripke's arguments largely rely on case intuitions meaning that at best he shows that languages containing non-rigid designators are unintuitive. Perhaps that's a reason not to adopt those languages. Moreover, even if Kripke's right and names have to be rigid designators, that's a semantic fact, not a logical one. It would still be the case that logical correctness is not entirely settled by the logical facts. See Ahmed (2007) for further discussion of Kripke's work.

Objection 2: the correctness of the inclusion or exclusion of non-rigid designators is determined by the logical facts. There are objective facts about identity. The identity predicate in a language needs to describe these facts. If the necessity of identity is true, then terms need to rigidly designate so that it does not. If necessity of identity is false, then non-rigid designators need to be allowed.

Reply 2: I'm sceptical of objective facts about identity, even granting the existence of metaphysical facts more broadly. Identity seems to have more to do with the co-reference of names (i.e. something semantic) than anything metaphysical. Identity is really just an object-language expression of the meta-linguistic facts of co-reference. But even granting that there are, languages with non-rigid designators can still successfully describe the facts about identity, provided that they also contain rigid designators. The rigid designating part of the language succeeds in capturing an important metaphysical structure and the non-rigid part serves the practical utility of having non-rigid designators.

Objection 3: Logical realists frequently conceded that it's possible to define languages in which certain objectively incorrect rules are valid or objectively incorrect rules are valid. They distinguish between logical rules being semantically and metaphysically correct (e.g. in Tahko (2021)). This is just another instance of that. You can define a language where **MSR** is semantically incorrect, but its metaphysical correctness is an objective fact.

Reply 3: The case is not like cases of mere semantic, but not metaphysical, correctness of a logical rule. The point is that what decides **MSR** is, in part, a fact about the nature of terms. There is no metaphysical fact that entirely determines **MSR**. There is no fact of the matter about **MSR** simply on the metaphysical level, because **MSR** depends on semantics facts as well as metaphysical ones.

In summary, then, at best, the logical facts do not resolve all logical disputes. Many contain semantic elements that are not resolved on entirely metaphysical grounds. This means that logical correctness is not entirely determined by the logical facts, even granting that such facts exist.

## 3.4 Why care about the logical facts at all?

This argument is an adaptation of an argument against ethical realism first presented in Korsgaard (1996) and more recently discussed in Peterson and Samuel (2021). Their argument goes as follows. Suppose that there are objective ethical facts. E.g. certain actions take on a special property if they are good and lack that property if they aren't. Why, ethically, should we care about this property? What makes this property really something worth having? Even if the realist can establish that there are so-called 'ethical' facts, they still have the task of establishing why those 'ethical' facts are actually relevant to ethics, why they're something that should guide our actions.

The ethical realist needs to establish two claims: (1) that there are objective ethical facts which can be described by ethical theories and (2) that these facts are relevant to the question of which ethical theory one ought use. Korsgaard (1996) and later Peterson and Samuel (2021) argue that even if (1) is true, (2) doesn't automatically follow.

The same argument can be made with respect to logical realism. The logical realist has to establish two claims: (1) that there are objective logical facts (**OBJ**) and (2) that these facts are relevant to logical correctness (**REL**). This section targets this second claim.

Suppose, using the example of §3.1, that fuzzy predication is objectively false. Properties in reality really are boolean and there are no partial facts about some object o having the property P. The use of any fuzzy predicate would be, strictly, incorrect. It would fail to describe the objective structure of reality. But this doesn't mean that we should therefore abandon the use of fuzzy predicates. They are incredibly useful in contexts where undue precision is unnecessary and either difficult or actively unhelpful to achieve.

Take something like the distinction between a hill and a mountain. Suppose that reality is Boolean and either (1) there is a crisp cut-off somewhere between hills and mountains or (2) there's no objectively tenable distinction between the two. All there is from objective reality's point of view are blocks of rock standing at particular heights.

I think it would be impractical to care too much about reality here. If (1) is true, that doesn't mean that the concept of "mountain" should be fixed to the objective boundary. The concept might be used more practically to indicate, say, the need for warmer clothes or to expect a longer hike. It might be intentionally misapplied to draw out certain important considerations. If someone primarily familiar with the Alps asks if England has mountains "sort of" is probably the best reply, to indicate a level of contextually appropriate nuance, irrespective of where the objective boundary between hills and mountains sits.

If (2) is true, logical realism might lead one to think that the concepts of "hill" and "mountain" should be abandoned, or at least arbitrarily precisified in a manner that serves no practical purpose beyond ensuring that our languages models the One True Logic. But that would clearly be impractical. Just because no objective boundary for vague terms exists doesn't mean that they don't serve practical use. The terms, for instance, might indicate the kinds of clothing or equipment one should bring when hiking. This is not something a reality's eye view perspective would necessarily care about, but it's certainly something we should.

The upshot is that, in this case, the inclusion or exclusion of a particular semantic object from the language (fuzzy predicates) should be a pragmatic matter; the metaphysics is only of limited relevance. But this has the upshot that the correct logic to use is also a pragmatic choice in some contexts not a strictly metaphysical one. Logic selection does not always care about metaphysics.

Moreover, there's nothing special about hills and mountains here, or even vagueness for that matter. The more general idea is that there are frequently cases where pragmatic concerns are more important than metaphysical ones.

Interestingly, some logical realists have been quite open to this kind of argument. They distinguish between the semantic and metaphysical correctness of a logical rule. Something like dialetheism, for instance, might offer the best account of the semantics of natural language liar sentences and hence be semantically correct for those languages, or be the most useful account of the liar and hence worth adopting, whilst nevertheless being metaphysically false (Again see Tahko (2021)). Moreover, the logical realists don't necessarily advocate for adopting languages whose logics are metaphysical correct. They might concede that objectively false logics are useful in some contexts. But these logics would nevertheless not be *correct*, merely a helpful falsehood.

Surprisingly, the Neo-Carnapian and the logical realists find a fair amount of agreement here. The Neo-Carnapian concedes to the Metaphysician that there are could be contexts where what the metaphysician calls 'logical facts' would matter (assuming such facts exist)[8]. When doing fundamental physics, for instance, this is might the case. Both agree that there are also contexts where those considerations take a back seat in favour of other kinds of consideration. They very plausibly agree on which contexts are which.

Where's the disagreement, then?

There are two important parts to Tahko's comments worth separating. First, the claim that there is a distinction to be made between semantic and metaphysical correctness. I don't think anyone's doubting that. Second, the claim that sometimes it's permissible to use objectively false but situationally helpful logic.

[8]Here my Neo-Carnapian breaks from Carnap quite sharply. Carnap could not acknowledge this as he believes metaphysical claims of this sort are framework-independent and hence pseudo questions. My Neo-Carnapian is not committed to Carnapian frameworks, though.

Logical realists could easily reject this second claim. One could do this outright, though that would be a little extreme. A more plausible route is to claim that in our most serious or important contexts, certainly any scientific context, it's unacceptable to use an objectively false logic. Though perhaps day-to-day contexts have more relaxed rules.

At this point, the Neo-Carnapian and the logical realist have essentially reached an agreement. The Neo-Carnapian is not claiming that there couldn't be certain logics that stand in some special relationship to certain logical facts, provided such facts exist. The Neo-Carnapian is also happy to admit that there might be contexts where the objective logical structure of reality, if such a thing exists, might be relevant for logic selection. Fundamental physics, for instance. They simply claim that this isn't uniform across all contexts. Any logical facts that do exist are only sometimes relevant to logic selection. The logical realist provides the correct account of logical correctness for some contexts, but only for some.

The logical realist might push back on this and insist on a difference between metaphysical logical correctness and normative logical correctness. A logic is metaphysically correct iff it correctly describes the structure of the logical facts (i.e. the logical realist's notion of correctness) but is normatively correct iff one may reason with this logic (in a given context). The logical realist might clarify that what they mean by correctness is metaphysical correctness. They were making no claims about what logics one can reason with.

If this is the case, then the only disagreement between the Neo-Carnapian and the logical realist is over what to apply the label "logical correctness" to. Both grant that there's a normative notion of correctness and the Neo-Carnapian. Now, this is clearly an entirely semantic dispute. It's merely a matter of where to apply the particular label "correctness". But nevertheless, the Neo-Carnapian's use of "correctness" has an advantage over the logical realist's: normative upshot. As discussed in §1.1, something I assume the concept of "correctness" is supposed to do is tell us which logical rules we may reason with. If the logical realist's notion of "metaphysical correctness" does not achieve this. The Neo-Carnapian's "normative correctness" does.

In summary, then, even if the logical realist can establish the existence of logical facts and show how they, in principle, determine a 'correct' logic, they are still left with the more challenging task of establishing why these metaphysical facts matter for the philosophy of logic and for logic selection. This secondary point, so this section argues, is where the metaphysical approach to the philosophy of logic falls short.

# 4 The One-Language Approach to the Logical Correctness

The One-Language Approach approach holds that:

**SEM:** Logical truths are semantic in that they are mind but not language-independent.

**OL:** The semantic truths, given some fixed language, are what matter for logical correctness. The interesting sense of logical correctness is correctness with respect to the semantic facts of some fixed language.

Amongst the major progenitors of this position is Dummett. He's certainly not the first or the only person within this context to advocate for this kind of view, but there's a relatively tangible intellectual history to be told that puts him in a central position within the last fifty years or so.

He says the following:

> *My contention is that all these metaphysical issues turn on questions about the correct meaning-theory for our language. We must not try to resolve the metaphysical questions first, and then construct a meaning-theory in the light of the answers. We should investigate how our language actually functions, and how we can construct a workable systematic description of how it functions; the answers to those questions will then determine the answers to the metaphysical ones.* (Dummett, 1991, p338)

Amongst these metaphysical issues are questions in the philosophy of logic. This can be broken apart into two claims:

1. Many metaphysical, including logical, issues *" turn on questions about the correct meaning-theory for our language"* - This is **SEM**.

2. *"We should investigate how our language actually functions, and how we can construct a workable systematic description of how it functions; the answers to those questions will then determine the answers to the metaphysical ones."* - This is **OL** where the fixed language is the actual language.

The later Wittgenstein arguably holds a similar position, Wright attributes Wittgensteinian influences to Dummett, but there are also ways of reading Wittgenstein's work that are more Carnapian[9].
Dummett's approach to the philosophy of logic (and theoretical philosophy generally) influenced his PhD student Crispin Wright. Whilst Wright and Dummett disagreed on a great deal, *Truth and Objectivity* (Wright, 1992) is in part Wright's reply to Dummett. The two agree on **OL**, though. In *Truth and Objectivity* Wright argues for a specific kind of Alethic pluralism. Crucially, Wright is interested in showing that *within the same language* it's possible to have multiple truth predicates[10]. On this basis, Wright does establish a kind of logical pluralism, but this isn't a major part of the work.

Wright's work went on to be a cornerstone of the modern alethic pluralism debate, with Lynch providing the next influential book. Probably the most notable trend since *Truth and Objectivity* is the shift towards domain-relative views. Wright argues for multiple truth predicates correct for the entire language. Most modern alethic pluralists argue for many truth predicates singularly correct for some domain within the language.

The step from any kind of alethic pluralism to a corresponding logical pluralism is not a complex one. Pedersen (2014) has presented the clearest outline of logical pluralism on the basis of a Wright-Lynch style alethic pluralism. Steinberger (2019) has criticised this view.

---

[9]If one interprets language games as, essentially, Carnapian frameworks, then the two start to look quite similar

[10]In Essay V: *Language, Truth & Logics* I call this 'narrow scope alethic pluralism', contrasted with the (Neo) Carnapian claim that there are many theories of truth correct for some language, the 'medium scope' claim

Crucially, though, throughout this entire debate, the alethic pluralists, and those like Pederson who wish to base their logical pluralism on their alethic pluralism, there is the assumption that monism or pluralism about both truth and logic needs to be defended within a given language.

They're not interested in showing that there can be many truth predicates, and subsequently many logics, true of one language.

There are more one-language theorists than merely those directly following Dummett via Wright. Most major versions of logical pluralism within the modern debate aim to show how logical pluralism is true within the same language, i.e. holding one's language fixed.

Beall and Restall (2005), for instance, are pluralists about the notion of entailment. They take there to be many different relations all of which are perfectly sensible entailment relations. There are many logics that are correct for some entailment relation. Again, this is all taking place within a given language. Beall and Restall show how logical pluralism could be true, given both a fixed language and without alethic pluralism.

Shapiro (2014) is another example. He takes logics to be models of truth-preserving inference within some language. He's a pluralist because he thinks there are multiple different, but all correct, ways of modelling truth within some language.

Lastly in Cook (2010)'s generally well-received overview of contemporary logical pluralism he explicitly endorses **OL**. He describes language-relative versions of logical pluralism as *"insubstantial"*, though gives no real argument as to why that is the case[11]. He outlines what he calls substantial logical pluralism which is explicitly logical pluralism, given some fixed language (and a fixed notion of the logical/non-logical divide). See also Cook (2023).

[11] Cook also rules out Varzi (2002)'s logical pluralism, a version of pluralism based on pluralism about the logical/non-logical divide, on the same grounds. This is puzzling given that he thinks Beall & Restall's pluralism is substantial and the two views are similar in the kind of logical pluralism they argue for. Beall & Restall's view, for reference, is pluralism about the notion of "all possible cases" in the definition of validity. Both argue that there's a particular part of the definition of validity that is ambiguous and this allows for logical pluralism on the basis of different precisifications of this ambiguity. One might go further and say that Varzi's pluralism about the logical/non-logical divide is a mechanism by which one could have pluralism about the range of possible cases. If this were the case, then Varzi's view would be a proper sub-view of Beall & Restall's.

I also wish to include the project of linguistic explication within this group. Unlike pure one-language views, explicators like Lewis (Lewis (1986) and Lewis (1998)) or Tarski (1931) are willing to revise their languages to some extent, but not much. Lewis, talking about philosophy generally, accepts that natural languages might not be consistent. The job of philosophy is to tidy natural languages up causing as little disruption as possible; to find the formally consistent nearest neighbour to our natural languages. Tarski (1931), at least on the common reading, sets out similar goals in relation to truth.

Whilst the explicators allow for some linguistic change so are not, in the strictest sense, One-Language theorists, a fixed language (typically our natural language) plays a very central role in what they're trying to achieve. Explicators have more in common with "pure" One-Language theorists than they do with Neo-Carnapians. My objection to One-Language theorists also applies to the explicators.

## 4.1 The Possibility of Linguistic Revision

The main argument against the one-language approach comes from the possibility of linguistic revision. Given that languages *can* be changed, the semantic rules of any particular language are not especially important in determining with which logic one should reason. Suppose, for instance, one's fixed language contains fuzzy predicates. It does not necessarily follow that one should reason with a fuzzy logic, as this might just mean that one should amend the fixed language (i.e. unfix it) in order to remove (or precisify) the fuzzy predicates.

Both the one-language theorist and the Neo-Carnapian accept **SEM**. Logical facts are determined by linguistic facts. They also agree that there is logical variation across languages. The one-language theorist does not deny that there are possible rival languages that could be constructed, modelling logics different from their own. What they deny is the *importance* of this for logical correctness.

One reason to think this connects back to the discussion in §1.1. Something logics are supposed to be able to do is provide us with norms for reasoning[12]. Ultimately, it's the validity of a rule in one's actual language, or the language one adopts within some more rigorous scientific context, that matters for the permissibility of reasoning by that rule. It is, after all, that language in which one reasons. Consequently, what matters for the philosophy of logic is logical correctness within the language one will actually be reasoning in.

The mistake here relates to the discussion of §2.2. It assumes that there are no norms relating to language selection relevant to the philosophy of logic. As discussed in §2.2, the Neo-Carnapian has a two-tier view of the norms of reasoning. There's an array of norms relating to language selection and then an array of conditional norms (hypothetical imperatives, to use a classical term) that state which inferences are permissible, given a particular choice of language.

That a particular inference is valid in some particular language is only sufficient for the conditional norm that if one uses that language, then one may reason by that rule. But the Neo-Carnapian's point is that one's choice of language is not forced and not beyond the scope of critique in the philosophy of logic. But this means that what's relevant for determining the rules one should reason by is not just a study of validity in one's own language (or some fixed language) but in a range of possible languages, along with a study of why one might choose one language over another.

As an example, suppose that it's the early 20*th* century and one's privileged language is the language of pre-quantum science. Suppose the correct logic for this language is classical logic. Suppose, then, that quantum phenomena are discovered and it turns out that a quantum logic would have certain beneficial features for describing this domain. This is obviously a contested point[13], but grant this for the sake of the example.

[12]This is a variation on the normativity argument against logical pluralism as discussed by Russell (2020) alone, Blake-Turner and Russell (2018). The version presented here makes the weaker claim that something logics should do is contribute towards the norms of reasoning. I don't claim that logics *are* normative theories of reasoning.

[13]see Putnam (1968) and Gibbins (1987) for an introduction to this topic.

Clearly what should happen is the adoption of quantum logic, at least within the confines of theoretical physics. But the one-language theorist can't endorse this move. They're interested only in the correct logic for their fixed language. They don't recognise the norms of linguistic revision, the kinds of norms that speak in favour of Quantum Logic within the example.

As a second example, suppose that English is the privileged language, that it has fuzzy predicates and that this is an undesirable result. For the one-language theorist, fuzzy logic is correct, and that's all there is to say of the matter. It is the correct logic of an English reasoner's language, so it is the logic that constrains what inferences they are allowed to make. The Neo-Carnapian sees another option. The reasoner might choose to change their language, moving into, say, some language English*. This is English, but with boundaries, perhaps arbitrarily, specified for each fuzzy predicate. By moving into English*, they are no longer bound by the logical rules correct for English. They are free to adopt whichever non-fuzzy logic is correct for English*.

Interestingly, the logical realist can also level a similar objection against the one-language theorist. Whilst they grant that one's language might fix the norms of reasoning in that language, that's entirely moot if the language doesn't model the objectively correct logic. Upon discovering, as in the Quantum case, that the logic of a language is false, one should change the language. The difference between the Neo-Carnapian and the Metaphysician here is simply that the Neo-Carnapian has a more expanded conception of the kinds of inter-linguistic norms that bear on language and logic selection.

This option fails because the one-language theorist fails to consider the possibility of linguistic, and hence logical, revision. It's not enough for the one-language theorist to simply show that a logical rule, in fact, does hold within some privileged language. They need to show, in addition, why this language shouldn't simply be changed.

Another option for the one-language theorist is to take a similar stance to the logical realist positions discussed in §3.4 which distinguishes between metaphysical and normative correctness. They can reject that when they talk about logical correctness, they mean anything to do with logic selection. The notion of "correctness" they have in mind is simply a relation between a logic and some fixed language.

They're not making any claim about how one should go about picking one's language or logic.

If the one-language theorist wishes to take this line, there's again some level of agreement between the one-language view and the Neo-Carnapian view. As with the discussion of metaphysical correctness in §3.4, a distinction might be drawn between correctness-for-a-given-language and normative correctness. The one-language theorist simply wishes to investigate how a given language can relate to a or many logics. Then they use the term "correctness" they mean correctness-for-a-given-language..

But the notion of correctness relevant to deciding how one ought reason is not correctness-for-a-given-language, given that linguistic change is possible. Thus, if the One-Language theorist makes this move, they all but concede that the Neo-Carnapian's notion, not their own, is the more important one for both practical and theoretical reasoning.

Note, as well, that all of this applies equally well to the explicator. The explicator is willing to change their language a little, but only for the purpose of ensuring internal consistency. The examples of reasons for linguistic change given all fall well outside that remit.

## 4.2 Reply: A change of subject?

There are a remarkable number of parallels between this debate and the debate between conceptual engineers and conceptual analysts. Conceptual analysts hold that when doing philosophical work on concepts, the aim is to understand our concepts *as they are*. The conceptual engineer, on the other hand, wishes to understand what they want particular concepts to do and to amend these concepts to better suit their purposes[14].

[14] Conceptual Explicators sit somewhere between the two. They allow for some degree of linguistic change, but only a minimal amount. In my view, this only gets the worst of both worlds. They are subject to the change of subject objection like the conceptual engineer but are lumbered with any consistent but undesirable features of their language like the conceptual analyst.

An analogue of the debate in §3.1 can take place between the conceptual analyst and the conceptual engineer. The conceptual engineer asks why they should care about their actual concepts. The conceptual analyst replies that it's because these are the concepts that they do, in fact, have. The conceptual engineer is unmoved, they are used to changing their concepts frequently so don't feel especially bound by the contingent features of their actual concepts, unless those features can be motivated as useful. A common reply by the conceptual analyst at this point is the so-called change of subject objection.

Consider, for instance, the question "Is being a stay-at-home parent a job?" Very plausibly, on the traditional meaning of the word "job" this is false. A job is paid economic labour, typically via an employment contract with a business. The production of goods or the provision of economic services in exchange for money. Stay-at-home parenting is, so say, not *economic* labour, in that it does not result in the production of goods or *economic* services, and is not typically paid.

For the conceptual analyst, that's all there is to say. For The conceptual engineer, however, this is just the beginning of the story. The concept "job", understood in a way that excludes domestic labour, has a meaning that encodes or promotes certain negative social values. It devalues the importance of domestic labour. They argue that, therefore, the concept of "job" should be changed to include stay-at-home parenting, and other forms of domestic labour.

A major reply to the conceptual engineering program is the change of subject objection. The change of subject objection holds that the conceptual engineer has done something invalid. They are no longer answering the question "Is being a stay-at-home parent a job?" but rather the question "Is being a stay-at-home parent a job*?", where "job*" is the concept job plus the conceptual engineer's modifications. The conceptual engineer hasn't changed the concept of "job", they've simply adopted a new concept job* and stuck the old label on it. They then haven't answered the question "Is being a stay-at-home parent a job?", they've just changed the subject.

This argument has interesting similarities to Quine (1970)'s and Restall (2002)'s meaning objection to Carnapian logical pluralism, an argument which also applies to my Neo-Carnapian view.

Here I present a slightly adjusted version of their argument to better fit the topic of linguistic revision. See my work elsewhere for a more faithful rendering of their argument and a more detailed and thorough reply.

Consider the question "Does $\phi \vdash \phi \vee \psi$?" Grant that, for our actual language, or for the privileged language, it does. The Neo-Carnapian is not moved by this fact alone. They could, for instance, adopt a three-valued language where disjunctive syllogism is false. What they're interested in is understanding the merits of adopting each of these languages over the other. But Quine and Restall object that they're simply changing the subject. The question was phrased in the original language, so it's a question about $\vee$, the disjunction of the original language. What the Neo-Carnapian does is try and answer the question "Does $\phi \vdash \phi \vee^* \psi$?" where $\vee^*$ is disjunction in their new language. They've simply changed the subject, not answered the question.

The conceptual engineer has a number of replies to the change of subject objection, three are presented here. See Belleri (2021) for a more detailed discussion of the various historic and prospective replies to the objection by conceptual engineers. Each reply has a natural counterpart for the Neo-Carnapian.

First, the conceptual engineer can argue that concepts can have their content changed without changing their numerical identity. In short, they argue that job=job*.

For the Neo-Carnapian to take this line, they would have to reject traditional truth-conditional semantics. I'm sympathetic to this option. I don't think truth-conditional semantics are appropriate when discussing languages with different notions of truth. However, that claim is more than a little controversial and I don't have a sufficiently worked-out alternative to present at present. I therefore won't advance this option at present.

Alternatively, the conceptual engineer can argue that whilst job and job* are not identical, *subjects* of discussion or inquiry are sufficiently coarse-grained as to allow for some level of conceptual change (Cappelen, 2018).

This is a potentially interesting approach and it would be interesting to see the details of this view born out in the specific context of logical inquiry. I won't develop this reply here, though.

The last option is to accept that there's a change of subject but argue that this is a dialectically permissible move. They accept that they're not answering the question "Is being a stay-at-home parent a job?", but that doesn't matter because it's an irrelevant question. It's a question formed with a faulty concept, one that they're abandoning, so it can be dismissed.

Analogously, the Neo-Carnapian could accept that they are answering the question "Does $\phi \vdash \phi \vee^{*} \psi$?" rather than "Does $\phi \vdash \phi \vee \psi$?", but they don't care. $\vee$ is a connective in a language that they've now abandoned. Rules about $\vee$ are no longer relevant for their reasoning.

To tie this back to the discussion of normativity, the Neo-Carnapian is interested in knowing what norms should govern their reasoning. Having established that they should reason in a language with $\vee^{*}$ rather than $\vee$, the conditional imperatives which follow $\vee$'s inference rules become irrelevant. This means that not only is the change of subject acceptable to the Neo-Carnapian, it's also required by their aims.

In summary, then, the Neo-Carnapian responds to the change of subject objection by accepting that there's a change of subject but denying that this is an issue. The change of subject is a feature, not a bug, of their view. It's baked into their two-tier conception of the norms of reasoning.

# Conclusion

This paper outlines three approaches to logical correctness: logical realism, the one-language approach and the Neo-Carnapian approach. It outlines the Neo-Carnapian view, contrasting it with Carnap's classic view (§2). It then explains the logical realist (§3) and one-language (§4) views, and gives an objection to each on behalf of the Neo-Carnapian view.

The Neo-Carnapian approach to correctness is ultimately preferred on the grounds that it is the only view with normative upshot. If the concept of "correctness" is supposed to entail something normative, i.e. that one may reason with a correct logic, only the Neo-Carnapian's view does this appropriately. The logical realist misses that there are contexts in which metaphysical considerations are not relevant.

The one-language theorist misses that linguistic change is possible.

## References

Ahmed, A. (2007). *Saul Kripke*. London: Bloomsbury Academic.

Armstrong, D. M. (1997). *A World of States of Affairs*. Cambridge University Press, Cambridge.

Ayer, A. J. (1936). *Language, Truth and Logic*. London: V. Gollancz.

Basu, R. and Schroeder, M. (2019). Doxastic wronging. In Kim, B. and McGrath, M., editors, *Pragmatic Encroachment in Epistemology*, pages 181–205. Routledge.

Beall, J. and Restall, G. (2005). *Logical Pluralism*. Oxford, England: Oxford University Press.

Belleri, D. (2021). Downplaying the change of subject objection to conceptual engineering. *Inquiry*.

Benacerraf, P. (1965). What numbers could not be. *Philosophical Review*, 74(1):47–73.

Benacerraf, P. (1973). Mathematical truth. *Journal of Philosophy*, 70(19):661–679.

Blake-Turner, C. and Russell, G. (2018). Logical pluralism without the normativity. *Synthese*, pages 1–19.

Burgess, J. and Rosen, G. (1997). A subject with no object: Strategies for nominalistic interpretation of mathematics. *Philosophical Quarterly*, 50(198):124–126.

Cappelen, H. (2018). *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.

Carnap, R. (1934). *Logische Syntax der Sprache*. Springer, Vienna. Translated into English 1937 by Amethe Smeaton under the title "Logical Syntax of Language", Routledge, London.

Cook, R. T. (2010). Let a thousand flowers bloom: A tour of logical pluralism. *Philosophy Compass*, 5(6):492–504.

Cook, R. T. (2023). Perspectival logical pluralism. *Res Philosophica*, 100(2):171–202.

Dummett, M. (1991). *The Logical Basis of Metaphysics*. Cambridge, Mass.: Harvard University Press.

Field, H. H. (2001). *Truth and the Absence of Fact*. New York: Oxford University Press.

Frege, G. (1879). *Begriffsschrift: Eine der Arithmetischen Nachgebildete Formelsprache des Reinen Denkens*. Halle a.d.S.: Louis Nebert.

Frege, G. (1893). *Grundgesetze der Arithmetik*. Hildesheim,: G.Olms.

Gibbins, P. (1987). *Particles and Paradoxes: The Limits of Quantum Logic*. New York: Cambridge University Press.

Griffiths, O. and Paseau, A. C. (2022). *One True Logic: A Monist Manifesto*. Oxford University Press.

Halbach, V. (2010). *The Logic Manual*. New York: Oxford University Press.

James, W. (1896). The will to believe. In Cahn, S. M., editor, *The Will to Believe: And Other Essays in Popular Philosophy*, pages 1–15. New York: Longmans, Green, and Co.

Kim, B. and McGrath, M. (2019). *Pragmatic Encroachment in Epistemology*. New York: Routledge.

Korsgaard, C. M. (1996). *The Sources of Normativity*. New York: Cambridge University Press.

Kripke, S. (1980). Naming and necessity. *Philosophy*, 56(217):431–433.

Lange, M. (2016). *Because Without Cause: Non-Causal Explanations in Science and Mathematics*. Oxford, England: Oxford University Press USA.

Leng, M. (2021). Models, structures, and the explanatory role of mathematics in empirical science. *Synthese*, 199(3-4):10415–10440.

Lewis, D. K. (1986). *On the Plurality of Worlds*. Malden, Mass.: Wiley-Blackwell.

Lewis, D. K. (1993). Mathematics is megethology. *Philosophia Mathematica*, 1(1):3–23.

Lewis, D. K. (1998). *Papers in Philosophical Logic*. New York: Cambridge University Press.

Liggins, D. (2006). Is there a good epistemological argument against platonism? *Analysis*, 66(2):135–141.

Maddy, P. (2007). *Second Philosophy: A Naturalistic Method*. New York: Oxford University Press.

Maddy, P. (2012). The philosophy of logic. *Bulletin of Symbolic Logic*, 18(4):481–504.

McSweeney, M. M. (2018). Logical realism and the metaphysics of logic. *Philosophy Compass*, 14(1):e12563.

McSweeney, M. M. (2019). Following logical realism where it leads. *Philosophical Studies*, 176(1):117–139.

Pedersen, N. J. L. L. (2014). Pluralism × 3: Truth, logic, metaphysics. *Erkenntnis*, 79(S2):259–277.

Peterson, C. and Samuel, J. (2021). The right and the wren. In Shoemaker, D., editor, *Oxford Studies in Agency and Responsibility 7*, pages 81–103. Oxford University Press.

Priest, G. (1987). *In Contradiction: A Study of the Transconsistent*. New York: Oxford University Press.

Priest, G. (2006). *Doubt Truth to Be a Liar*. New York: Oxford University Press.

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

Priest, G. (2010). *Logic: A Brief Insight*. Sterling.

Putnam, H. (1968). Is logic empirical? *Boston Studies in the Philosophy of Science*, 5.

Quine, W. V. (1981). *Theories and Things*. Cambridge, Mass.: Harvard University Press.

Quine, W. V. O. (1960). Word and object. *Les Etudes Philosophiques*, 17(2):278–279.

Quine, W. V. O. (1970). *Philosophy of Logic*. Cambridge, Mass.: Harvard University Press.

Restall, G. (2002). Carnap's tolerance, meaning, and logical pluralism. *Journal of Philosophy*, 99(8):426–443.

Russell, G. (2020). Logic isn't normative. *Inquiry: An Interdisciplinary Journal of Philosophy*, 63(3-4):371–388.

Shapiro, S. (2014). *Varieties of Logic*. Oxford and New York: Oxford University Press.

Steinberger, F. (2019). Logical pluralism and logical normativity. *Philosophers' Imprint*, 19.

Tahko, T. (2014). The metaphysical interpretation of logical truth. In Rush, P., editor, *The Metaphysics of Logic: Logical Realism, Logical Anti-Realism and All Things In Between*, pages 233–248. Cambridge: Cambridge University Press.

Tahko, T. E. (2021). A survey of logical realism. *Synthese*, 198(5):4775–4790.

Tarski, A. (1931). The concept of truth in formalized languages. In Tarski, A., editor, *Logic, Semantics, Metamathematics*, pages 152–278. Oxford University Press.

Tarski, A. (1943). The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, 4(3):341–376.

Varzi, A. C. (2002). On logical relativity. *Philosophical Issues*, 12(1):197–219.

Wittgenstein, L. (1922). *Tractatus Logico-Philosophicus*. Routledge & Kegan Paul.

Wright, C. (1992). *Truth and Objectivity*. Cambridge, Mass.: Harvard University Press.