Change of Logic and the Adoption Problem

MSc Thesis *(Afstudeerscriptie)*

written by

Horia Lixandru

under the supervision of **Dr. Luca Incurvati** and **Dr. Giorgio Sbardolini**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense:	**Members of the Thesis Committee:**
3rd June 2025	Dr. Aybüke Özgün (Chair)
	Dr. Luca Incurvati
	Dr. Giorgio Sbardolini
	Dr. Karolina Krzyżanowska
	Dr. Thomas Schindler

Institute for Logic, Language and Computation

Abstract

The Adoption Problem is the claim that certain basic logical principles cannot be adopted, meaning that if one does not already infer in accordance with them then one cannot start following them by explicitly being told the principle. I argue that the Adoption Problem rests on misleading assumptions. The right diagnosis is that someone has failed to understand the principle due to not possessing the concepts used to state the principle. One can acquire a logical concept by being immersed in an activity in which the concepts are embedded. In the case of quantifiers, I will follow an idea due to Hintikka, and show how one can learn the meaning of the quantifiers by learning to play a kind of game called the 'seeking-and-finding' game.

Keywords: Adoption Problem, Universal Instantiation, Concept learning, Hintikka

Acknowledgements

First and foremost, I want to thank my supervisors, Luca and Giorgio. Their guidance helped me make this thesis what it is today, and I couldn't have asked for a better supervision team.

Many thanks to Karolina Krzyżanowska and Thomas Schindler for agreeing to be my examiners, for their insightful questions during the defense, and for all the discussions during my MoL years. Thanks to Jeroen Smid and Simone Picenni for their feedback on earlier drafts of Chapter 4, and to Sabina Domínguez Parrado for comments on Chapter 5. I also want to thank Corine Besson for sending me a draft of her upcoming book, and for taking an interest in my project and my ideas! Thanks also to the audiences at The Inaugural Symposium on Logic in the Arab World and The Dutch Research School of Philosophy (OZSW) 12th Annual Conference 2025 at TU Delft, for comments on presentations about early results of my thesis.

Many thanks to my friends, both within the MoL and beyond it. To my dear friends in Amsterdam, the illogicians: Frank, Laura, Radu, Joris, Andrea, Zarah and Antoine; as well as Sarah, Martijn, Lamia, and Luca, who have taught me that when I feel it's too dark, I should first stop closing my eyes. This thesis was written while listening to a lot of music, and I want to give a special thanks to Linkin Park (who came back just as I was beginning this thesis) and Sigur Rós, for fuelling my writing sessions. Most of my writing was done in the MoL room, so I want to also thank all the great people I have met in this time for creating this great community.

Finally, thanks to my parents, for their unwavering support in all I have ever done. This thesis is for them.

Contents

CHAPTER 1

Introduction

There is no mathematical substitute for philosophy

Saul Kripke

An animal is a cat if it is a domesticated member of family Felidae. Something is a number if it conforms to Hume's Principle, and a logical connective is a conditional if it conforms to rules such as Modus Ponens. How did you learn what a cat is? Presumably when you were a kid you were pointing at some animal and asked an adult *What is that?* How did you learn what a number is? Presumably you started the long process of counting, and by means of some cognitive processes, you learned to associate the sound 'one' to the number One.[1] How did you learn to use conditionals? Presumably by hearing a lot of conditional statements in natural language, imitating them, and observing their effects on conversations.

In all of the above examples, there is a discrepancy between what something is, and how we come to learn that something is what it is. No one learned what a cat is through zoology, we observe animals with certain features and call them cats. No one learned numbers through Hume's Principle, we learn how to count and numbers somehow conform to those axioms. Finally, we were not told explicit rules for doing logic, we use logic and see that we accept these rules when we are told them.

Our understanding of certain concepts without certain practices would be, in some sense, deficient. One could learn what a cat is entirely by following

[1] Following convention, I will denote concepts by small caps throughout this thesis.

zoological definitions, but if they've never seen a cat, intuitively we think that there is something missing in their knowledge. Likewise for someone learning numbers exclusively through Hume's Principle but without having ever learned to count.

A recent line of thought, originating in Birman (2024), Kripke (2024), and Padro (2015), proposes a novel problem for the philosophy of logic, arguing that basic logical principles cannot be explicitly adopted. In other words, suppose you tell someone who has never used Universal Instantiation that *Universal statements imply each of their instances*. In order to start using the rule, they will need to apply Universal Instantiation since the rule itself is a universal statement. This conclusion is supposed to confer a special epistemic status to certain logical principles, with some of the ones mentioned so far being Universal Instantiation, Modus Ponens (Birman 2024; Finn 2019, 2021; Kripke 2024), Adjunction (Susanszky 2023), and the law of non-contradiction (Kripke 2024).

This thesis will run counter to this line of reasoning. I will argue that the Adoption Problem is misleading, and that one of the assumptions used to generate it fails. The way I will do so is by claiming that someone who fails to perform easy instances of a basic logical rule has failed to grasp the concept(s) used in formulating the rule. For example, someone who is not disposed to say that from *Everything is self-identical*, *Frank is self-identical* follows has failed to grasp the concept All. Then, acquiring a concept is not something that is done via rules for using a concept, but is a process that is much different. For example, children do not acquire their knowledge of numbers via formal principles, such as Hume's Principle, but do so via some sort of activity, for example counting. Combining these two insights, I will conclude that the Adoption Problem, if understood at the level of rules, is a false problem, or an uninteresting one, and if formulated at the level of concepts, then it lacks bite from the start.

Here's how this thesis will proceed. In the following chapter, I will give an in-depth overview of the Adoption Problem, the argument used to reach it, and its supposed bearing on other philosophical problems. Chapter 3 will be a critical overview of some recently proposed solutions to the Adoption Problem, and I will argue that most of the proposed solutions fall short of being satisfactory. Chapter 4 will argue that someone who fails to perform easy instances of a rule actually lacks the concept used to formulate the rule, and as such it should not be surprising that they have failed to adopt it. The final chapter will be the positive contribution of this thesis. I will sketch and defend an account for how someone could acquire the required concept by presenting an insight due to Hintikka (1973). Chapter 6 will conclude and reflect on possible takeaways of my argument.

CHAPTER 2

The Adoption Problem and the Epistemology of Logic

The 'Adoption Problem' is based on an until recently unpublished series of lectures given by Kripke in the '70s, now partially published as Kripke (2024). The name 'Adoption Problem' is due to Padro[1] (2015) and, until recently, has received very little attention (for the very few exceptions, see Berger 2011; Stairs 2006, 2016). Kripke claims that logical rules cannot be adopted and that intuitive reasoning cannot be dislodged by some formal theory. The issue Kripke raises comes in the context of Putnam's (1968) argument for revising logic on empirical grounds, more specifically adopting quantum logic in order to accommodate reasoning with quantum phenomena. Setting the waning popularity of quantum logic aside, the issue raised is nonetheless a common and pervasive one in the philosophy of logic, namely (rational) revision of logic. If (some) logical principles cannot be adopted, this might give those unadoptable logical principles a special epistemic status, and spell trouble for those who take logic to be revisable (for a small sample, see Priest 2006; Quine 1951; Warren 2020).

I claim that Kripke's argument, and in particular Padro's rendition of the argument, is unsound. I will show how the tension in the thought-experiment of Harry, who supposedly fails to adopt a logical rule, can be solved. I will reject the assumption that Harry understands the rule, but keep the idea that rules are constitutive for the meaning of certain expressions. Birman (2024) suggests that this will raise the Adoption Problem at the level of concepts,

[1] Note that Padro and Birman are the same person, so I will refer to her by either, depending if I am refering to her PhD thesis or to her recent paper on the problem.

but I will show how this can also be solved. A key move will be denying that adopting rules is how we come to acquire certain concepts. In other words, concept acquisition is a substantially different process from adopting a logical rule, with a key analogy being the fact that children do not acquire number concepts by learning Hume's Principle. My diagnosis will be that Harry fails to adopt in virtue of not understanding the rule, which is due to him lacking the concepts used to formulate the rule. Thus, it is not the rules that are at fault, but Padro's suggested method of rule-learning.

In this chapter I will walk through the Adoption Problem as presented by Kripke and Padro, lay down the argument explicitly, and give a brief motivation for my solution. I will then compare the Adoption Problem to Lewis Carroll's (1895) Tortoise, as well as highlight the importance of the Adoption Problem for other broader debates in the epistemology of logic.

2.1 The Adoption Problem

In this chapter, I will unpack the assumptions behind the Adoption Problem, present the argument given in order to reach it, and place it in a broader philosophical context. To start, I will give a brief overview and offer a small sample of the issues at stake.

Let's start right away with the statement of the problem. This is the *Adoption Problem*:

> (AP) : certain basic logical principles cannot be adopted because, if a subject already infers in accordance with them, no adoption is needed, and if the subject does not infer in accordance with them, no adoption is possible. (Padro 2015, pp. 41-2; cf. Birman 2024, p. 39)

One exegetic note before we move on. We can distinguish at least two versions of the (AP). One is Kripke's (AP), and the other one is Padro/Birman's (AP). In terms of strength, Kripke's version is stronger. For now, it suffices to say that Padro's version is concerned with *basic* logical principles, while Kripke argues that adopting *any* logical principle is impossible. Padro's version allows for adopting more complex logical principles if basic ones are already in place. Moving forward, I will focus primarily on Padro's version of the argument.

In the course of this section I will unpack the problem and walk through the reasons Padro gives for arriving at this conclusion. The first half of the statement seems to me uncontroversial. If I already infer using some rule, maybe the most that one could do is make the rule *explicit* (i.e., state the rule explicitly, say that it is valid, and that I have been using it all along, and if I do indeed infer in accordance with it, then no issue should arise on my part). As such, it is the second half of the definition that raises problems.

To illustrate this second horn of the (AP), the following thought-experiment is usually given. Suppose we meet Harry, who (1) has never performed an inference in accordance with *Universal Instantiation* (UI), and (2) has never encountered the (UI) principle beforehand:

$$(\text{UI})\ \frac{\forall x Px}{Pa}$$

This inference rule says, simply, that if *All things are P*, then, for any individual, that individual is also *P*. For example, from *Everything is self-identical*, via (UI), I can deduce that *Frank is self-identical*, or from *Everything is black* that *The Munttoren is black*. Now, let's see if we can get Harry to perform inferences using (UI).

Suppose we tell Harry the (UI) rule, stated in English as follows:

Universal statements imply each of their instances

Let us also assume that he trusts us, and that he accepts the principle as valid. Subsequently, we want to check Harry's understanding. Suppose we tell him that *Everything is black* is true, and ask whether *The Munttoren is black* is true. His reply is expected: he doesn't know. In order for him to be able to perform the inference, he would first have to recognise *Everything is black* as a particular instance of a universal statement. In other words, he has to be able to infer according to (UI) which, by assumption, he cannot. Thus, us telling Harry the rule explicitly was of no help. He is in no better position than before. An analogous situation arises for someone who does not reason in accordance with *Modus Ponens* (MP). If we meet such a person and we tell them the (MP) rule explicitly, then they will not be able to perform an (MP) inference, despite explicitly accepting the rule.

In the rest of this section, I will present in more detail the Harry argument in favour of the (AP), and lay down the argument in a more straightforward fashion, laying the groundwork for the rest of this thesis. But before that, I want to give more detail as to what *adopting* a logic means, since Padro (2015) gives the term a technical sense. This will be the focus of the next subsection.

2.1.1 *What does it mean to adopt a logic?*

So far, we have worked with a more-or-less intuitive understanding of what *adopt* means. To both Kripke's and Padro's distress, such uses of *adopt*, which according to them are commonplace in the philosophical literature, are unclear, or worse, incoherent:

> And people have felt that we have a well-defined notion of 'adopting a logic'. In fact, even now lots of the literature supposes that we have some standard procedure for 'adopting a logic' (Kripke 2024, p. 9).

It will be helpful to have in mind what exactly *adopting* a logic amounts to. Padro's preferred objects of adoption are rules, so we will move forward with rules in mind. Briefly, what is a rule? In a nutshell, it is an imperative of the form *If C, do X*.[2] Without aiming to go further in this discussion than an elementary exposition, the following two features seem to capture, at least partially, our understanding of a rule (Reiland 2024):

- Our actions can *accord* or not to a certain rule.[3] For example, relative to the rule *If you're tired, go to sleep!*, if I go to sleep when I'm tired I perform *in accordance* with the rule. Note that it is irrelevant whether I *explicitly* follow it. I can go to sleep when I'm tired without having to think about the aforementioned rule.
- Rules can *guide* our actions. For example, if I endorse the rule *If you're tired, go to sleep!*, then when I'm tired, me going to sleep can be (partially) explained by my endorsement of the rule. Likewise, I can consider the rule when I'm tired in order to decide what to do.

Padro (2015, p. 42) uses *adopt* as a technical term. We can say that to *adopt* a rule R:[4]

Adopting a Rule

(1) One must not yet reason according to R
(2) One comes to accept R as a *good* rule
(3) It is in virtue of this acceptance that one comes to rationally reason in accordance to R

I take condition (1) to be straightforward and not in need of much discussion. If we already reasoned in accordance with the rule, then perhaps we don't need to adopt it, but rather make it explicit or something similar. In any case, *adoption* seems, uncontroversially, to be about going from one practice to another, or from one rule to another, or from no rule to some rule. Therefore, I don't think anyone would dispute condition (1).

Likewise, condition (2) seems fairly uncontroversial, but perhaps less so than (1). Why would adoption require us to accept a rule as good? The

[2] Rules can also be framed as oughts (*If C, you ought to X*), as permissions (*If C, you may X*), or in other similar forms. I will go with the imperative form for two simple reasons. The first, and the most important, I do not think the (AP) rests on a particular view of rules. Harry's problems are, I think, insensitive to the formulation of rules, since the issue consists in applying the rule itself, which will arise on any particular version. The second reason is continuity, as the imperative form of rules is the one favoured by Kripke (1982).

[3] Moving forward, I will stick to this basic view of rules, ignoring many important issues. For example, I am not considering what are usually called *weak imperatives* (von Fintel and Iatridou 2017, §13.2), such as *Sure, go ahead*. Sticking to the intuitive model of rules mentioned before will be enough for the (AP).

[4] One might complain that logical rules are not given in the form imperatives, and are merely descriptive. I deal with this issue in the following subsection where I'll present the argument for the (AP).

distinction that I think comes here is one between adoption and imitation. Imitation does not require me to recognise a rule as good, only to grasp its content and act in accordance with it. Thus, what makes me act in accordance with the rule is not something about the rule, but a reason that could be thought of as external to it (e.g., wanting to fit in a social group, etc.).

Condition (3) is required to distinguish *adopting* a rule from merely acting *in accordance* with one. If it just so happens that I never cross the street on a red light, I will act in accordance with traffic rules without explicitly following them. There is no consciousness of the rule involved in my acting in accordance with it. In cases of adoption we want someone to use the rule because they accept it, and not due to irrelevant reasons such as luck.

Now that we have an overview of the basic statement of the (AP), the thought-experiment motivating it, as well as a precise sense of what adoption means, we are ready to move on to the argument Padro (2015) makes based on the Harry-scenario. This is the focus of the next subsection.

2.1.2 *The Argument for the Adoption Problem*

In this subsection, I will walk through the argument given in favour of the (AP) in more detail. I will first highlight the core issue that gives rise to the (AP), and then present the argument given to support the conclusion that the core issue licenses the (AP).

First, let us walk through a more clear rendition of the Harry story. More succinctly, Harry has to follow the steps below in order to perform the required inference:

(2.1)	$\forall x Px \vDash Pa$	(UI)
(2.2)	Everything is black	Premise
(2.3)	Everything is black $\vDash$ The Munttoren is black	Application of (UI)
(2.4)	The Munttoren is black	Application of (UI)

We want Harry to follow this pattern because we want him to reason to the conclusion (2.4) *as a consequence* of accepting (2.1), and thus we want to rule out other ways of reaching the conclusion (e.g., already believing the conclusion for other reasons). This is the way we can summarise Padro's (2015) presentation of the Harry case. Steps (2.1) and (2.2) act as premisses given to Harry, and we assume he accepts them. The crucial step that leads to the (AP) is (2.3). The idea that being able to apply (UI) before one can start using (UI) rests on the assumption that (2.3) is itself an application of (UI). More precisely, abusing logical notation, Harry has to apply (UI) to the following:

(2.5) $(\forall x Px \vDash Pa) \models$ (Everything is black $\vDash$ The Munttoren is black)

Thus, we see that, on Padro's account, in order for Harry to be able to reach the conclusion in (2.4), he would first need to be able to apply (UI) in (2.3) in order to recognise the inference as a particular instance of a universal statement. Harry first has to instantiate the (UI) schema for the case of the predicate 'is black'. In other words, he has to apply (UI) to go from (2.1) to (2.3). As this issue will come back later, I will label it **Presupposed**, and assume that, if true, it implies that an agent cannot use a rule unless it is already part of their practice:

> **Presupposed:** Initial application of a basic logical rule presupposes knowledge of the rule.

We can view **Presupposed** as the practical equivalent of what has been dubbed epistemic *rule-circularity* (Boghossian 2000; Dogramaci 2010; Dummett 1978) in the literature on justifying basic logical principles. Rule-circularity means employing (a meta-logical analogue of) the logical rule in justifying the logical rule, such as when proofs of the validity of Modus Ponens involve using Modus Ponens in the metalanguage. In the practical case, performing inferences in accordance with (UI) requires competence in using (UI), at least according to proponents of the (AP).

Now, let us lay down more clearly the assumptions and the argument that seem to lead to the (AP). Some of the assumptions that will be relevant for the discussion have been named in boldface. Padro's argument is intended as a reductio of the thesis labeled below **Inferential Cognitivism**. I will first give the argument in full, and then spend the rest of this section discussing it.

Adoption Argument

(1) **Inferential Cognitivism:** Performing inferences according with R requires (possibly tacit) acceptance of R
(2) The Harry Scenario: Suppose there is an agent H who does not use a basic logical rule R. H is explicitly told the rule R, H comes to assent to R as a good rule, i.e., they accept R, H accepts the premiss(es) of R, but H does not infer that the conclusion licensed by R follows.
(3) **Understand:** H understands R, meaning that the agent understands the words used to formulate the rule and how they come together.
(4) **Presupposed:** Initial application of a basic logical rule presupposes knowledge of the rule
(5) **Immediate:** Adopting a basic logical rule is an immediate process.
(6) H does not use R despite understanding R, assenting to R, and accepting the premisses of R

(7) H cannot adopt R by explicitly being told R
(8) ∴ An agent cannot adopt a basic logical rule by explicitly representing the rule
(9) ∴ (1) is false

Before we proceed, let's set some technicalities aside. I will assume that (UI) is a generally valid[5] principle, and that Harry is otherwise a competent English speaker. For the sake of simplicity, I will gloss over differences between natural language words associated with universality, e.g., *all*, *every*, *everyone*, *everything*, and use simply the term *all* moving forward. This simplification might be warranted on linguistic grounds, since the concept ALL seems constitutive of all other expressions of universality (Wierzbicka 1996, p. 46f). Next, what I mean by *explicitly told* is that the agent gets the rule from an external source, e.g., told by someone, read in a book, divine inspiration, etc. This excludes the possibility of Harry evolving the practice some other way, e.g., by being exposed to the environment or being conditioned (assuming he would be capable of undergoing such a process). Finally, I take as granted that Harry *should* infer the conclusion when asked whether it follows from the accepted premisses, thus setting aside any issues related to normative aspects.

The focus of the literature on the (AP), as well as this thesis, is centred on the bolded assumptions: **Understand, Presupposed, Immediate**. To block Padro's Adoption Argument, one or more of the given steps must be blocked, and as mentioned, the focus has been on **Understand, Presupposed, Immediate**, and in this thesis I will argue that we should reject **Understand**. For reasons of space, I will not discuss possible relationships between these assumptions, and instead focus on each one individually. Having set aside some of the technicalities, now I am going to motivate more in-depth the more contentious assumptions. Let's start with **Immediate**.

Birman assumes that adoption is something that needs to happen immediately, or very fast. She explicitly excludes the possibility of Harry being trained or conditioned to reason in accordance with (UI):

> First, mere causal connections, where the relation between Harry's acceptance of the principle and his inferring in accordance with it are fortuitous, do not count as cases of 'adoption'. The principle should be used to generate his UI inferential practice. Second, the possibility of developing the practice by training or psychological conditioning is excluded. There should be no doubt that Harry's acceptance of the principle is the reason and explanation for his inferring in accordance with it (Birman 2024, p. 40).

[5]Thus, we exclude the possibility of Harry being a free logician.

The condition for a successful adoption of a basic logical rule is that the agent, in virtue of grasping the rule, is able to apply it, and is able to do so immediately. Acceptance, in this sense, leads to unmediated adoption.

One might complain, as Devitt and Roberts (2024) do, that learning a skill takes time. Essentially, I might know all the steps, but transferring all this explicit knowledge into performance is nonetheless a matter of training. For example, I might know how to play individual chords, and know a chord progression in theory (even better, I might know *a lot* of music theory), but I could still be incapable of playing the guitar. Transforming this theoretical knowledge into practice is a matter of training, and we could treat logical knowledge in the same way.

I think we can overcome this issue in the following manner. We can make a distinction between supposed *basic* rules, and more complex ones. I do not mean *basic* in the sense of primitive rules of a formal system, but rather basic in the sense of feasability of inferring. Borrowing an example from Cherniak (1986, p. 30), inferring $\neg q \rightarrow \neg p$ from $p \rightarrow q$ might be easier than inferring $\exists x \forall y(Fx \rightarrow Gy)$ from $\forall xFx \rightarrow \forall yGy$. To push the idea further, reusing the guitar analogy, the basic skill is simply putting your fingers on the fret in such a way as to play a chord, while the more complex skill is playing a chord progression, or a song, keeping rhythm, etc. We can likewise use an example from logic. We can think of the rules of a Natural Deduction system as corresponding to basic skills, and being able to put together proofs as a complex skill.[6] Basic logical rules are unlike skill-based rules or heuristics because they can be adopted without training. We can thus still assume that **Immediate** holds for such basic rules of inference. While perhaps adopting these rules might not be enough to properly pick-up a practice (as perhaps adopting all the rules of English might not be enough to start speaking English), the most basic of steps can be immediately applied, even if one might not know how to connect the steps.

Next, a very quick remark regarding **Presupposed**. Arguably, it seems that, on pain of regress, *some* logic is required in order to apply or consider applying any kind of inferential rule. Indeed, this is the main criticism of Quine's empiricist epistemology of logic, which will be the subject of §3.1. However, here the focus is different. Padro argues that we require applying *the very same* rule. Applying (UI) for the first time requires already knowing how to use (UI), and likewise applying (MP) for the first time requires already knowing how to use (MP). I will postpone more in-depth discussions of this assumption until §3.5. This brief overview will suffice until then.

[6] One can also cash out the distinction in terms of *constitutive rules*, distinguishing them from heuristics, or rules-for-playing-well, or a skill-determined rule, if there is such a thing. A constitutive rule is, taking an example from chess, the rule that the knight moves in an L-shape, whereas a skill-determined rule might be *If you can capture a pawn without risking your knight, capture the pawn*.

There is also the minor assumption that knowledge of logical principles is manifested by acting in accordance with these principles. Since Harry does not manifest such logical knowledge, this gives grounds to believe that he also lacks the requisite logical knowledge. Of course, this might not be a knock-down argument, since the Tortoise in Carroll's story also refuses to manifest logical knowledge, i.e., to infer, but it seems clear enough that the Tortoise knows the rule of inference (as it knows exactly what Achilles needs to do, but adopts a sceptical attitude). Likewise, I will take this assumption on board for now, and discuss it more in-depth once I present my own solution to the (AP). The rest of this section is devoted to a short motivation for rejecting what is, in my opinion, one of the least-disputed assumptions of the Adoption Argument: **Understand**, which I will restate for the scenario at hand:

> **Understand:** Harry understands the (UI) rule, meaning that he understands the words used and how they come together.

We assume that Harry understands what the rule says, but is clueless as to how he should apply it. Thus, to reach Padro's conclusion, Harry must understand all the components of the rule, universal quantifier included. If one found out that Harry didn't understand the words used to express the rule, then it will be no wonder that he isn't able to apply the rule. How can he apply something he doesn't understand in the first place?

However, blaming Harry for not immediately understanding and applying a rule he has never seen, nor used, seems akin to blaming a child for not being able to use numbers after telling them Hume's Principle.[7] In order for a child to learn the numbers, a radically different approach is needed, and indeed actual learning proceeds differently (Buijsman 2017; Carey 2009a). Simply stating formal rules is a highly inefficient and unnatural method of learning.[8] This is the approach I will follow in this paper. I claim that Harry's inability to follow the rule is due to him not understanding the rule. This is my initial claim:

Claim

Harry is unable to adopt (UI) because he does not understand the rule.

[7]The principle stating that the number of *F*s is identical to the number of *G*s if and only if *F* and *G* are equinumerous. For more on Hume's Principle and its relationship to neo-logicist mathematics, see the essays in Hale and Wright (2001, Part III). In both cases, there is nothing at fault with the principle, but with the agent. I will return to the analogy between Harry and Hume's Principle in §4.4 and §5.1 below.

[8]A comparison can be drawn between Plato's *Meno* and what I will present in this paper regarding the learnability of concepts, as well as innateness. I only mention this as a possible interesting historical connection to the issue at hand. However, such a comparison is beyond the scope of this paper. Thanks to Karolina Krzyżanowska for suggesting this comparison.

In other words, it is not the rules that are special, but it is the agent that is lacking something. This will be the novel contribution of my thesis, and of course the claim will be further refined in the course of my argument. As a preview of the next pages, I will use the rest of this chapter to discuss the (AP) more in depth, and the next chapter to discuss currently available solutions and argue why they should be rejected. My main task in Chapter 4 will be to show why rejecting **Understand** is a principled move. To that end, I will show why Birman (2024) thinks rejecting **Understand** is not a good solution, and show how we can resist her move.

2.2 What Carroll said to Kripke

As mentioned, the (AP) is based on Kripke's interpretation of Lewis Carroll's short note on Achilles and the Tortoise. Carroll (1895) seems to question the possibility of ever drawing an inference according to a logical rule: Modus Ponens (MP). It has been dubbed the *paradox of inference* (Pavese 2022), as the conclusion seems to conflict with our intuitions. We seem to draw inferences all the time. The case of the Tortoise in Carroll's dialogue seems similar to Harry's situation. In both cases an inference is not drawn, despite the acceptance of a logical principle. In this section, I will present (without aiming to preserve the literary quality of Carroll's prose) the main ideas behind the dialogue between Achilles and the Tortoise, with an eye towards Kripke's interpretation of the note.

Suppose we ourselves meet the pesky (and by now, very old) Tortoise. He has kept its skeptical attitude, and it is our turn to continue Achilles' labour. Suppose we try to get the Tortoise to perform the following inference:

(1) If I'm in Amsterdam, then it's rainy.
(2) I'm in Amsterdam
(■) ∴ It's rainy

The Tortoise, as expected, refuses to see the necessity of the conclusion, and claims that he must also accept:

(3) If (1) and (2) are true, then (■) is true

Then, the following, we believe, might suffice:

(1) If I'm in Amsterdam, then it's rainy.
(2) I'm in Amsterdam
(3) If (1) and (2) are true, then (■) is true
(■) ∴ It's rainy

We think, wrongly once more, that this is enough and that the Tortoise must finally accept (■). "And why *must* I?" (Carroll 1895, p. 279). Fortunately, unlike Achilles, we know how this story goes.[9] Unfortunately, like Achilles, we seem to be in no better position to answer the Tortoise's demands.

[9]If not: an infinite regress ensues. The Tortoise demands that, before he can accept (■), he has to accept (4) If (1) and (2) and (3) are true, then (■) is true, and similarly for an implication

Many morals have been drawn from Carroll's (1895) short dialogue.[10] Here I will instead focus on the one essential to Kripke's discussion: the importance of differentiating between logical principles and rules of inference.[11]

Very roughly, while logical principles are theorems of a certain (meta-)logic, e.g., the law of excluded middle ($\forall p(p \vee \neg p)$), the law of non-contradiction ($\forall p \neg(p \wedge \neg p)$), and so on, *rules of inference* are tools that allow us to go from true premisses to a true conclusion, e.g., Modus Ponens, and are usually thought to belong to the meta-language[12]. Then, the Tortoise is refusing to apply a rule of inference, and instead demands more and more logical principles to be granted as extra premisses of the putative inference. However, as Priest puts it succintly, "no extra entailments, taken as premises, can force a person to infer i.e. to *do* something" (Priest 1979, p. 291, my emphasis).[13]

Let me make the difference clearer through an example. Taking once again the canonical case of (MP), the rule of inference is usually spelled out as follows:

$$\text{(MP)}\ \frac{p \to q \qquad p}{q}$$

However, the logical principle equivalent to this rule is the following:

(MP principle) $\qquad p \wedge (p \to q) \to q$[14]

The distinction can also be further emphasized intuitively as follows: While we can say that the (MP principle) is true, saying the same about (MP) would be a category mistake. Inference rules do not seem to be truth-bearers, but are instead usually said to *preserve-truth*, i.e., whenever all the premisses are true, the conclusion is true as well.[15]

With the above in mind, let's render the Tortoise's argument, in a similar fashion as to how I presented Padro's argument in the previous section.

---

(5), (6) . . . , *ad infinitum*. The Tortoise will keep doubting the necessity of the conclusion and demand additional premises.

[10] For a much more thorough analysis and discussion, see Besson (2018).

[11] This distinction was one of the first proposed to solve Carroll's problem, see for example Cook Wilson (1926, pp. 443ff), B. Russell (2010, Ch. 3).

[12] Or, at the very least, they are somehow outside the argument itself: "invoking a supporting rule is external to the deductive structure of an argument, while adding a premise is internal to" (Smiley 1995, p. 731).

[13] The view put forward by Priest, while not unintuitive, might not be universally endorsed. For example, in ethics, if we understand that, say *helping others is good*, then we have a *pro tanto reason* to help others (Incurvati forthcoming). In other words, we might think that validity is intrinsically normative, and if one fails to grasp this, then one has failed to grasp the concept.

[14] As I presented them here, both the principle and the inference rule are schematic. If one so desires, we can formulate the MP principle as a propositionally quantified statement instead of a schematic one. Pavese (2022, p. 6), for example, takes this approach. I will discuss the difference between the two in §3.5.

[15] Hanna (2006, p. 57) makes a similar remark.

Tortoise Argument

1. The Tortoise understands (MP)
2. The Tortoise assents that (MP) is a good rule of inference
3. **Rational-If** If an agent understands (MP), assents to (MP), and accepts the premisses of an (MP) inference, then they know they should use (MP)
4. **Tortoise-demand** The *strengthened* version of an argument must be accepted before the argument can be accepted.
5. The Tortoise knows that they should use (MP) but they refuse to do so.

---

6. ∴ The Tortoise is akratic.

I am borrowing the term *strengthened* version of an argument from Thomson (2010) to capture the demands of the Tortoise: Before they can conclude that from p and $p \to q$ they can deduce q, they must include $(p \wedge p \to q) \to q$ as an additional premiss, and so on. More precisely, we call the *strengthened* version of an argument $\varphi_1, \ldots, \varphi_n \vDash \psi$ the argument $\varphi_1, \ldots, \varphi_n, (\varphi_1 \wedge \cdots \wedge \varphi_n) \to \psi \vDash \psi$. The **Tortoise-demand** leads to the infinite regress of premisses, since before I can accept an argument, I must accept its strengthened version first, but to do that I must accept the strengthened version of the strengthened version... and so on. While this assumption might seem similar to **Presupposed** from Harry's scenario, they are nonetheless different. In Harry's case, **Presupposed** generates no regress, but simply blocks the adoption of the principle.

Now, what is the relationship between Harry and the Tortoise? The Harry scenario raises one question regarding the classical solution to Carroll's paradox: How can we make someone *adopt* a logical rule? Achilles tries to make the Tortoise adopt a logical rule, but phrases it as an extra logical principle. In Harry's case, we formulate the rule explicitly. However, the problem seems to be that we cannot get someone to adopt a rule by giving them the rule:

> If someone cannot see the inference from p and if p then q, to q, you will not help him by telling him: 'Look, the following is correct: if p and if p then q are true, then q is true'. This is because if you just add this on as an extra premise, one will still have to use Modus Ponens to get anywhere (Kripke 2024, p. 18).

While, as mentioned, the distinction between principles and rules is only one important lesson from Carroll's note, it is the one Kripke (2024) draws upon in arguing for what has been subsequently dubbed the *Adoption Problem*. Indeed, the echo of Carroll's lesson, as well as of Priest's remark above are clear:

> [O]ne cannot adopt the logical laws as hypotheses and draw the consequences. You need logic in order to draw these consequences [...] One has to just think, not in terms of some formal set of postulates but intuitively, that is, one has to *reason*. (p. 20).

Let's take stock. There are both similarities, but also important differences between Harry and the Tortoise. The most important one is that the Tortoise can be said to understand the principle, but refuses to apply it due to a skeptical attitude, whereas Harry has issue with *applying* the given principle.[16] The summary is rendered below:

Harry vs. Tortoise

Adoption Harry:	Carroll's Tortoise:
■ Does not know the rule beforehand ■ Has a trusting attitude, accepts the rule ■ Moral: Accepting a rule does not enable one to use the rule	■ Knows the rule beforehand ■ Has a skeptical attitude, demands extra premisses before following the rule ■ Moral: Rules of inference cannot be extra premisses

With these in mind, I now turn to another issue, namely why should we care about the (AP) at all? I will briefly go through some of the main implications the (AP) might have, if true, on broader debates in the epistemology of logic.

2.3 The Importance of the Adoption Problem

Having gone through the main statement and argument in favour of the (AP), it might seem as if I have glossed over one important aspect. What is the relevance of the (AP)? Why should it matter if some basic logical principles turn out to be unadoptable? In this short subsection, I will highlight some potential connections between the (AP) and contemporary philosophy of logic.

The Background Logic Problem

One of the assumptions, **Presupposed**, seems to require some kind of logical competence before one can actually use any logical theory. More bluntly, evaluating logic requires logic. This is usually called the *Background logic* problem (Hlobil 2020; Martin 2021; Wright 1986).

[16] Padro (2015, §2.2.1) provides a more in-depth exegesis of Kripke's interpretaion of Carroll's dialogue.

Kripke's claims are similar to others made in the literature, both before and after his seminar. This is akin to Shapiro's (2000) 'super-logic', the logic we use to evaluate any other logic, or Hanna's (2006) 'proto-logic', a very weak form of logic which is presupposed by all and every system of logic imaginable, which is itself *a priori* and unrevisable. This problem has also been historically dubbed the *logocentric predicament*, following Sheffer (1926, p. 226), the claim that "If the system constitutes the universal logical language, then there can be no external standpoint from which one may view and discuss the system. Metasystematic considerations are illegitimate rather than simply undesirable" (Goldfarb 1979, p. 353).

Conventionalism about logic

Supposing the (AP) is true, we cannot adopt basic logical principles, and furthermore, some logical principles are not, in a sense, up for grabs. This runs counter both to historical forms of conventionalism (Carnap 1937), and to contemporary conventionalism in the philosophy of logic (Warren 2020; Woods 2023).

Conventionalism about logic states, simply, that logical truths are true in virtue of our conventions, be they explicit or implicit. Thus, the logical truths are fully explained by our conventions, and further, we are free to choose the conventions we want. This is the *principle of tolerance* of Carnap (1937), the idea that "in logic, there are no morals" (p. 52). The (AP), if true, shows this cannot be the case, with the argument running similarly to Quine's (1936) original argument against conventionalism. If we cannot adopt certain principles, then we cannot impose 'a logic' upon us.

The Justification of Basic Logical Laws

Finally, the (AP) is supposed to bear on the problem of justifying our basic logical laws. Usually, the conundrum is that proving the soundness of some logical rule uses the very same rule in the proof, and so other methods are needed. According to Birman (2024, pp. 57ff) the (AP) bears on most approaches to justifying logical laws (for an opiniated overview, see G. Russell 2015), such as those based on intuition (Priest 2016; Wilkinson 2025; Wright 2004a), meaning-constitutiveness (Boghossian 2000, 2003), or empirical verification (Quine 1951; Williamson 2017).

CHAPTER 3

Current Perspectives on the Adoption Problem

In the previous chapter I have introduced the Adoption Problem (AP), and walked through the argument Padro gives for it, explicitly stating the argument and the required assumptions. I have compared Carroll's dialogue to Kripke's interpretation, and highlighted some important consequences of the (AP). In this chapter, I will discuss proposed solutions to the (AP), and argue that each one fails to overcome Padro's adoption argument, or that Padro's argument can be adapted to respond to the given challenges. I will begin by discussing Quine's empirical conception of logic.

3.1 Quine and anti-exceptionalism about logic

The most discussed problem in the epistemology of logic is the *justification* problem, i.e., how we justify our logical principles and rules. One way to do this is by declaring logic to be no different from science, and proceed to justify our knowledge of logic in the same way we justify our knowledge in other fields of inquiry. This is, historically, Quine's position, which has seen a recent surge in popularity under the name *anti-exceptionalism about logic* (Hjortland 2017) and is currently, arguably, the most popular position in the epistemology of logic.

Kripke (2024) directly criticises Quine's epistemology of logic and Putnam's proposal which is directly influenced by Quine. In this section, I aim to give an overview of Quine's epistemology of logic, focusing on his views

presented in *Two Dogmas of Empiricism*, as that is the version on which Kripke is focused upon.

Quine's view can be summarised in a fairly straightforward manner: everything is a hypothesis in a holistic web-of-belief, and everything is, in principle, open to revision. Quine himself makes this very explicit:

> Re-evaluation of some statements entails re-evaluation of others, because of their logical interconnections – the *logical laws being in turn simply certain further statements of the system, certain further elements of the field*. [...] Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of the excluded middle has been proposed as a means of simplifying quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle? (Quine 1951, pp. 296-7, my emphasis)

Quine, on the face of it, seems to blur the aforementioned distinction between principles and rules. After all, everything is a hypothesis in the web-of-belief. Thus, we must assume that logical principles and rules of inference must both be phrased as further hypotheses of our web-of-belief. Dummett proposed a similar reading:

> Quine's thesis [that nothing is immune to revision] involves, however, that the principles governing deductive connections themselves form part of the total theory which, as a whole, confronts experience. Presumably, in order to avoid Achilles and the Tortoise troubles of the Lewis Carroll type, we must recognise the total theory as comprising rules of inference as well as logical laws in the form of valid schemata and their instances: but there is no reason to suppose that Quine draws a distinction between the status of such rules as against laws like Excluded Middle; they too must be equally liable to rejection under a heavy enough impact from without (Dummett 1973, p. 596).

Going further, I want to focus on two interrelated replies that have been raised against Quine's view and which are related to the (AP). The first one I will call the *framework problem*, and the second one the already mentioned *background logic* problem, taking inspiration from the laws-of-thought debate in the philosophy of logic (see the Introduction to Miguens 2020), and from the anti-exceptionalist literature.

The framework problem can be phrased as follows: 'Some principles and rules of inference cannot be part of the web-of-belief since they govern the very function of the web and the rules for revision.' What I mean by that is that certain principles seem to be required in order for the web-of-belief to function as intended, and thus must be in a sense *external* to it. For example, the process of revision begins once an inconsistency is detected in our belief web. However, how are we able to detect inconsistencies if we do not endorse the law of non-contradiction,[1] or something similar to it (Arnold and Shapiro 2007)? If I have $\{\varphi, \neg\varphi\}$ as a subset of my belief-web, I would need (the schema or the instances) of the law-of-contradiction as well in order for my belief to be deemed inconsistent and to trigger a revision of my belief web. Likewise, other revision principles, such as Woods' (2019) *logical partisanhood* or Quine's own *maxim of minimum mutilation*, seem to also have this framework-status:

> Logic is in principle no less open to revision than quantum mechanics or the theory of relativity [...] If revisions are seldom proposed that cut so deep as to touch logic, there is a clear enough reason for that: the principle of minimum mutilation. The maxim suffices to explain the air of necessity that attaches to logical and mathematical truth (Quine 1986, p. 100).

The framework problem is similar to the background logic problem, but I think it's important to distinguish between the two, as their focus is different. The framework problem says that laws of logic are simply the laws that frame thought, i.e., what meaningful thought is. The background logic, on the other hand, mentions that in order to be able to draw the consequences of some logical theory, one needs *some* logic in order to do so. I agree that these two problems are intimately connected, but I will treat them separatetly for the aforementioned reason.

That such principles should instead be considered the working-principles of the framework,[2] instead of further nodes in the web-of-belief, seems to be in line with Carroll's lesson. Certain principles cannot be taken as further premisses, or hypotheses, since this would trigger a regress or worse (from the Quinean point of view), leave some statements immune to rational revision.

The second problem is the problem of the background logic. As mentioned, this is closely connected to the previously described framework problem, but the focus is different. The background logic is the one Kripke explicitly brings out (although he does not use this term, his way of putting the problem indicates the background logic problem):

[1] More exactly, the law of non-contradiction *and* the principle of explosion, $\varphi, \neg\varphi \vDash \bot$, and also some form of deductive closure (i.e., if I believe p and $p \to q$ (or $p \vDash q$), then I should believe q).

[2] Pavese (2021b), not directly attacking Quine, makes a similar distinction between input and structure in the context of Carroll's regress. Premisses for an inference are the input, while the inner workings of reasoning are the structure that enable inferences to even take place.

> [*I*]*f* the logical truths are mere hypotheses, just like anything else, and one can 'adopt' them as one will, how, unless one has a logic in advance, can one possibly deduce anything from them? (Kripke 2024, p. 15)

This latches on Quine's remark above that "Having re-evaluated one statement we must re-evaluate some others, whether they be statements logically connected with the first or whether they be the statements of logical connections themselves" (Quine 1951, p. 297). Kripke is thus questioning how exactly are we to determine whether some statements are *logically connected* if we don't have a logic to tell us. Shapiro (2000) also makes a similar point:

> Since nothing is outside the web of belief, the Quinean would have it that the identification of the correct logic is part of the web. In particular, logic itself is subject to modification the way anything in the web is. [...] Which logic do we use to assess the consequences of different logics? Is there a correct logic for that, and is this super-logic also just a bunch of nodes in the current web? Regress threatens. Is the super-logic analytic, a priori, or incorrigible? (Shapiro 2000, p. 338)

More briefly, suppose we want to see how a revision from Γ to Γ' influences other beliefs. since the consequence relation is also part of Γ, which consequence relation should we use in order to determine whether $\Gamma' \models \varphi$ or not? Should we use the old consequence relation, $\models_{\Gamma}$, or the new one $\models_{\Gamma'}$? Or, as the last part of Shapiro's quote mentions, is there some 'super-logic' which governs belief revision but which itself is not open to revision?

By means of summary, Quine's position, and that of the anti-exceptionalist more generally, is nicely given as part of Hjortland's programmatic article:

> Logic isn't special. Its theories are continuous with science; its method continuous with scientific method. Logic isn't a priori, nor are its truths analytic truths. Logical theories are revisable, and if they are revised, they are revised on the same grounds as scientific theories (Hjortland 2017, p. 632).

Based on the assumptions mentioned in the previous chapter, it should be clear that Quine and the anti-exceptionalists reject **Presupposed**. Since logic is no different from science, the laws of logic have no special status and we are free to revise as needed. On Quine's picture, "it is *because* we chose to accept [the logical principles] as part of our web of belief that we are committed to them and their consequences" (Padro 2015, p. 109). This is exactly what Kripke and Birman argue against. Some logical principles, such as (UI), are unlike scientific hypotheses in that they cannot be freely adopted. They must already be in-place in order for any evaluation to take place. If

we understand 'logic' as theory, of course we are able to revise which logical theory is the correct one, but the interpretation at stake is 'logic' taken as inferential practice, and *that* is the one that has a special status.[3] We cannot simply choose a theory and 'adopt' it, some logic has to already be in place for that to work. Using Quine's own metaphor, "in his own conception we wouldn't have a web to begin with" (Padro 2015, p. 116).

In the rest of this section, I aim to critically review some of the solutions proposed in the literature which explicitly engage with Birman (2024) and Kripke (2024). I will summarise each one individually, and highlight why I find them unsatisfactory, but also what I think they get right. Importantly, I believe that combining each individual insight will allow for a proper solution, which will be the subject of the last chapter of this thesis. In particular, the solution I propose, while offering a different initial diagnosis, takes some points from some of the solutions proposed. For the purpose of this section, I will sometimes make reference to my claim that Harry does not understand the rule, even if my argument towards that conclusion will be the subject of the next chapter. I believe this will serve to better highlight both my criticism, and better understand my proposal in Chapter 5 below.

3.2 A BEHAVIOURIST REPLY

Devitt and Roberts (2024) argue that the definition of the (AP) is too restrictive. They suggest instead that we can come to reason in accordance with certain principles by conditioning, or by simple training. What they have in mind is an analogy to the way skills are learned in the real world. They take issue Padro's exclusion of this possibility of solving the (AP):

> [T]he possibility of developing the practice by training or psychological conditioning is excluded. There should be no doubt that Harry's acceptance of the principle is the reason and explanation for his inferring in accordance with it (Birman 2024, p. 40).

Explicitly learning skills is a top-down process (Devitt and Roberts 2024, p. 67). That means that we first have propositional knowledge, i.e., explicit representation of rules (e.g., *In order to* φ, *do* $p, q, r, \ldots$). Then, by training, this knowledge becomes automatic. The declarative knowledge acts as a starting point for learning a skill, and once the skill is learned, this declarative knowledge might well be superseded by practical knowledge.

Adopting a logical rule follows the same principle on their account (Devitt and Roberts 2024, p. 68). Harry could choose to undergo training in order to reason according to (UI) in virtue of him recognising (UI) as a good principle. After considering and rejecting a *self-training* solution (Devitt and Roberts

[3] I will say more on the difference between logical theory and logical practice in §3.4 below

2024, §6-7),[4] they suggest Harry to hire a logic coach. This logic coach is supposed to tell Harry *when* to apply (UI), something that Harry can't do by himself, and repeat the process with multiple such examples. Harry should, after much practice, become disposed to infer in accordance with (UI). This disposition is rational, since Harry accepted (UI) as a good rule, and voluntarily decided to undergo training in order to infer according to the rule. This acceptance is the same kind at play in Birman's original scenario; Harry simply trusts us.

Devitt and Roberts reject **Immediate**. According to their solution, we should not expect that theoretical knowledge is immediately transferable into practical knowledge. For example, some athlete might have good theoretical knowledge of how a difficult move is performed, but that won't mean that they will also be able to perform it simply based on understanding those rules. The same expectation should apply in Harry's case. He simply needs to be trained in order to perform in accordance with (UI).

One problem I find with Devitt and Roberts (2024) is that their proposal rests on the existence of such a logic coach.[5] Even if we grant that their solution might solve the (AP) for Harry, we can raise the following question: *Who* is training the logic coach? If the logic coach is also trained by another logic coach, we can keep asking the same question. If some logic coach, however, has not undergone such training, but instead adopted (UI) without any coaching, the (AP) can be raised for that coach. Thus, despite the fact that I believe their argument to have some force when it comes to allowing training as a possible solution to the (AP), the way they implement it seems simply to kick the (AP) down the chain of learning, and I do not see how the training view can be salvaged by appealing to a sort of oracle without raising the aforementioned issue.

However, regress issue aside, I think there is a more pressing worry regarding their proposed solution. I believe the way Harry is being trained falls prey to a Chinese-Room type argument (Searle 1980). To quickly go through the original argument, suppose you are locked in a room, and someone gives you papers written in Chinese, which you do not speak, and you are supposed to return an output. In the room there is an extensive list of instructions for how to match the given characters to the required output. The person outside the room might think that you speak Chinese, but you are only following a set of instructions, without really understanding Chinese. The logic-coach answer might not fully fit the pattern, but the worry of understanding still comes up. The training Harry undergoes is mere conditioning, meaning that Harry will come to associate universal statements with their instances, but

[4] The rejection of this solution is based on the fact that Harry might fail to recognise the *conditions* of applying the rule. Therefore, he could not learn on his own to reason in accordance with (UI) and (MP), since recognising the proper conditions presupposes ability to apply (UI) and (MP).

[5] Hattiangadi (2023, p. 71) also points out this flaw in their argument.

one could worry if this would really count as *inferring*. In cases of association, we simply have a causal-temporal connection between the assertion of two statements, but inference is thought to be more than mere transition in thought. Most accounts try to substantiate the intuition by placing certain restrictions on inference. For example, Boghossian (2014) suggests what he calls the *Taking condition*, that in an inference the agent *takes* the premisses to support the conclusion. One could very well argue against this specific condition, but most accounts of inference place some similar restriction in order to differentiate the act of inferring from simple transitions in thought.[6] Thus, we are able to form many dispositions by conditioning, but would we want to call them all acts of inferring? Something seems to be missing, and I think Birman is right to reject the proposal of Devitt and Roberts as a genuine case of adoption. Harry will not use (UI) *because* he accepts it, but he will only be trained to do so.

Nonetheless, I think the spirit of their reply, namely the inclusion of training, has some appeal. In the context of the (AP), we assume that a logical rule is something that can be followed immediately. If I tell someone the (UI) rule, then in virtue of its simplicity I will assume that they will need almost no time in order to follow it. However, the background in which the rule is given, namely that of language and concepts, *is* something which requires extensive training in order to acquire. While someone might want to resist the idea that someone having the concepts will require training, someone who does *not* possess the concepts used to phrase the rule will need training, not in order to follow the rule, but in order to acquire the concept.

3.3 A COGNITIVIST REPLY

Chudnoff (forthcoming) suggests, like Devitt and Roberts (2024), that the (AP) can be overcome. Unlike Devitt and Roberts, he thinks the right approach is to reject **Presupposed**, which he calls the *Inference Assumption*:

> Inference Assumption: Generating a representation, e.g., an inferential seeming, of what general inference rules mandate in a particular case requires making an inference (p. 7).

For Harry to be able to apply the (UI) rule, he must first *apply* the general rule for the particular case, which seems to require yet another inference, i.e., the one of recognizing the particular instance as an instance of a universal statement.

[6] Just for a sample, see Blake-Turner (2022), Hlobil (2019), Marcus (2020), Peacocke (2024), and Valaris (2014).

How should we understand the idea of an inferential seeming? The following example should give some intuition:[7]

$$\text{(AMS)}\ \frac{\text{If I'm in Amsterdam, then it's rainy} \qquad \text{I'm in Amsterdam}}{\text{It's rainy}}$$

$$\text{(AMS?)}\ \frac{\text{If if I'm in Amsterdam, then it's rainy, then I'm in Amsterdam}}{\text{I'm in Amsterdam}}$$

The (AMS) example should seem valid to anyone who sees it, and should be immediately so. The (AMS?) one should be less evident, even though it is valid by Peirce's law. Then, an inferential seeming seems to be an intuitive feeling of validity, that we have in virtue of e.g., being in the habit of inferring according to such laws, or grasping some rule, etc., and should trigger in the case of (AMS), but not in the case of (AMS?) (Chudnoff 2024).

Our basic cognitive architecture looks like it is following rules that *seem* like (UI) and (MP), but "the appearance is merely an artifact of this particular, and at best heuristic, way of representing productions" (p. 9). In other words, we merely represent the workings of our cognitive architecture *as if* they are following (UI) and (MP), mainly because we are used to representing rules as generalized conditionals, but that does not mean that our cognitive architecture *actually* follows these rules. Then, what is actually happening is that our cognitive architecture is *matching* patterns, which are not truth-apt in the same way as the elements of an implication are, to actions, with these systems being labelled *production* systems. These systems are "simple mechanisms" (ibid.). What it is doing is not necessarily something that *we* are doing.

The proposal to solve the (AP) runs as follows. Firstly, notice that "[t]he correct production must be learned, and typically this occurs through instruction and example" (p. 10). Thus, supposing someone grasps (UI) and stores it as a rule in their memory, the cognitive architecture will work as follows (p. 11):

> If the goal is to evaluate an inference, And the premise is 'For all x, A' and the conclusion is 'A(t/x)' for some A and t, *Then* retrieve an inference rule for that kind of inference.
>
> If a retrieved inference rule classifies an evaluated inference as correct *Then* generate an inferential seeming of as of its conclusion following from its premises

Then, supposing we want to evaluate the inference

[7] Indeed, as I will mention later on, it should give intuition to anyone who is *not* like Harry, i.e., someone who has a grasp of (UI) and (MP).

$$\text{(Self-identical)}\ \frac{\text{Everything is self-identical}}{\text{Frank is self-identical}}$$

Our production will first *match* the pattern of (Self-identical) to the form of (UI) (and matching is assumed to *not* be an inference). Then, an inferential seeming is generated, leading one to experience (Self-identical) as correct, without the cognitive architecture itself employing (UI) or (MP). Then, *if* you act on this seeming, you are *inferring* by following (UI). The crux then is that when you infer, you follow (UI), but the generation of the inferential seeming does not require (UI).

However, I am not fully convinced by Chudnoff's reply. The whole proposal, in my opinion, rests on whether or not these inferential seemings do the work they are supposed to do, and indeed if in Harry's case these seemings do actually kick in. For example, couldn't Harry be in the same position regarding (UI) as we are regarding (AMS?)? In other words, since (UI) is something he has never performed, why will he obtain such an inferential seeming? To push this idea further, why do we have an inferential seeming for (AMS) but nor for (AMS?)? I believe this is due to the centrality that a rule such as (MP) has with respect to the concept IMPLICATION. The way we currently use the concept IMPLICATION makes it so that we are immediately in a position to recognise (AMS) as an instance of (MP). Likewise, we are in a position to recognise (UI) because we possess the concept ALL. As will become apparent in the next chapter, I believe that Harry lacks the concept ALL, which is why he cannot perform inferences in accordance with (UI). Thus, I agree with Chudnoff that the generation of an inferential seeming is not an inferential step, but I think that only someone who already possesses the concepts, or who undergoes conditioning, will obtain that seeming. For example, if I am in an environment where everyone affirms the consequent, I will easily obtain an inferential seeming for affirming the consequent, but not for (MP).

Nonetheless, as with the previous solution due to Devitt (2006), I believe Chudnoff also gets something right. He goes one step further, and argues that the training has some goal, and that something happens at the cognitive level for Harry. However, I am not sure what kind of "instruction and example" (Chudnoff forthcoming, p. 10) will be important for Harry. If it the same one as for Devitt and Roberts, then the same issues apply. Thus, if no inferential seemings appear, no adoption is possible on this proposal. It is not fully clear if Harry has these inferential seemings at all, since the examples simply mention that he has not performed an inference according to (UI). It is not mentioned whether such seemings are nonetheless present. However, as mentioned, I think that Chudnoff gets right the idea that a more involved cognitive process must happen in order for Harry to 'adopt' (UI). The solution in the next section follows a similar idea, that of training, but embeds it in a larger context, that of a *practice*.

3.4 a pragmatist reply

Finley (2024) suggests that both Carroll's regress, as well as the (AP) are

> an attempt to undermine inquiry into logic itself, whether that be by arguing that inquiry of the right type (e.g., neutral/impartial inquiry) is impossible, or by arguing that inquiry into logic itself is somehow limited by logic's role in reasoning (p. 9).

Pushing this idea further, he mentions that

> In the case of the adoption problem, part of the intent is to pose a challenge to rational revision of logic, which presumably involves a form of inquiry into logic itself in terms of reasons for or against revising or adopting some formal theory of logic. If one cannot adopt basic logical rule, then inquiring into possible revisions seems useless since they could not be enacted (p. 10).

The (AP) is understood here as being connected to the *background logic* problem mentioned before. The idea is exactly the one that Kripke (2024) argues for, namely that we cannot choose a formal system of logic *independent* of some form of reasoning, which is understood to have some sort of implicit logical structure. The choice of logic will be directly affected by our reasoning, or by the background logic. Then, the worry is that this background logic is not an impartial arbiter of theories:

> [T]he logic employed to derive the consequences of theories is bound to favor theories based on that logic over theories based on competing logics (Field 1998, p. 12).[8]

The pragmatist reply tries to dislodge the apparent fundamentality of logic presupposed both by Kripke and those that raise the Background Logic problem.

> On the suggested pragmatist position, we never need to adopt or even think about basic logical rules in the *docens* sense before inferring in accordance with some pattern in the *utens* sense (Finley 2024, p. 22).

Then, Finley's solution also proposes a rejection of **Presupposed**. Before I move on, let me explain exactly what Finley means by "logical rules in the

[8] Williamson (2013) makes a similar point in the context of choosing metaphysical theories. Both Field and Williamson question the supposed *neutral arbitrer* view of logic.

docens sense". A common distinction[9] found in the literature is that between logic-as-theory and logic-as-practice. The jargon usually attributed to them is that of *logica docens* (logic-as-theory) and *logica utens* (logic-as-practice). In other words, *logica docens* is the logic as formal theory of validity, e.g., logic as found in logic textbooks, and is usually (but doesn't have to be) spelled out in a formal way, while *logica utens* is logic understood as theory of (actual) reasoning.[10]

Granting this distinction, the interesting question is how these two notions are related. Subsequently, this distinction is also useful in order to identify exactly in which sense Kripke takes logic to be immune to adoption. This subsection will thus map some positions between logical theory and logical practice, with an eye towards the (AP).

The most radical position is that logical theory has nothing to do with logical practice. If we understand logical practice as simply put *reasoning*,[11], this thesis amounts to the fact that reasoning is not guided by any formal theory of logic. The most known proponent of this thesis is Harman (1986).

A second position, the opposite of the aforementioned one, is one where we instead identify reasoning with following a logical theory. This is the view advanced by Perice, who writes:

> We never can really reason without entertaining a logical theory (Peirce 1931, §4.476)

The (AP) seems to position Kripke in direct opposition to Harman, but not on the same grounds as Peirce. Both Kripke and Harman place the practice of reasoning above logical theory, but arrive at different conclusions (Kripke that logic *is* reasoning, Harman that logic is irrelevant for reasoning). However, I think they would both agree on the fact that our inferential practices have priority over any formal logical theory. Then, whether we call what we do when we reason *logic* or not, that seems to be a matter of terminology in this case.

How does this bear on the Harry scenario? We can assume that Harry has accepted (UI) into his logical theory, meaning that, at least *prima facie*, he has accepted a formal principle. What the story is supposed to show is

[9]This distinction can also be found under various names. For example, Williamson (2024, p. 418) distinguishes between *scientific* logic and *folk* logic, corresponding to *logica docens* and *logica utens*, respectively, while Resnik (2004, p. 179) talks of (lower-case) *logic* and Logic, the former being the *logica utens*.

[10]The jargon in the contemporary literature seems to be attributed to Priest (2014), but dates back to the medieval era, and is also common in the writings of Charles Peirce, see Peirce (1931, §2.204)

[11]In fact, it is hard to see how one might argue that reasoning and logical practice, understood as *logica utens*, come apart. If we understand the logical practice broadly as our way of reasoning, then the identiy is clear by definition. If we understand reasoning as more broadly, e.g., including abduction, supposition, imagination, etc., then logical practice is seen as part of these activities.

that he failed to change his logical *practice*. In other words, logical practice is not something that can be changed by a change in logical theory, and if the (AP) is true, then logical practice comes conceptually before the logical theory. However, nothing prevents changes in logical theory. With this overview in sight, let's return to Finley's argument.

Finley's comment, the one given before discussing the *logica docens*, seems in accordance with what Birman (2024), Kripke (2024), and Padro (2015) have to say about the (AP). Our logical practice is conceptually prior to any logical theory on both their views and the pragmatist view. The idea that we can revise our logical theory is not an issue on both accounts. On Kripke's account, this is viable so long as we do not undermine intuitive reasoning (e.g., the revision of Aristotelian logic to account for empty terms). On the pragmatist account, we use logical theories to reflect on our conduct, and we can revise these theories in order to adopt new norms or resolutions, as logic for the pragmatist is a normative discipline, i.e., it sets norms of reasoning.

The point of contention is the revision of *logica utens*. Here the pragmatist view of Finley appeals to what he calls "the tribunal of experience and their optimism in the power of inquiry" (p. 23). There is explicit reliance on some prior inferential practice:

> if experience can lead to some foothold in a *logica utens*, then formal rules can be adopted and one can use them to further improve one's reasoning habits and dispositions (p. 23).

This inferential practice, based on habits, nature, or some similar experience, is presupposed when it comes to adopting new rules. What we do then is use a logical theory in order to modify our habits and dispositions (ibid.).

The weak point of Finley's argument seems to me the reliance on prior experience for the inferential practice. What do we do when an agent does not have the prior experience, like Harry? Finley seems to dismiss out of hand cases like Harry:

> I am optimistic that experience has led everyone's *logica utens* (by nature or nurture) to have sufficient starting places to some inferential behavior that is in accordance with general rules and could lead to adoption (p. 23, fn. 73).

The question is what could the pragmatist say in cases where no prior experience is available, as in the case of Harry? This seems hard to answer, since the case of the pragmatist seems to rest on some prior inferential practice that could act as a launching pad for further reflection and adoption of logical rules. Without such a prior practice, the pragmatist seems unable to help Harry. I am overall sympathetic to the pragmatist approach to logic, but I believe that Finley's optimism, while not misplaced when it comes to real

people, thought-experiments like Harry might not be satisfiably captured by such an approach.

Then, I believe the argument falls short of a satisfying solution. Supposing indeed experience is needed, the question then is *what* kind of experience? It is true, as a matter of fact, that we use quantifiers in our everyday language. However, Harry does not, so what is missing? What sort of experience requires being able to infer in accordance with (UI), or to use quantification? This is where I think my solution improves upon the one suggested by Finley. As mentioned, I am overall sympathetic to Finley's approach, but in order for his argument to go through, a stronger line of reasoning is needed. Then, Finley's argument rests on the assumption that someone like Harry will have "sufficient starting places", but what would those starting places be? In order to help someone like Harry who doesn't seem to have had such a starting place, it might do well to nudge him in that direction, but evidently it is hard to know which way to push if we don't know what we are aiming at.

3.5 Pattern-Recognition and Schemas

Another solution which rejects **Presupposed** is one proposed by Besson (2019). Recall that Birman claims that the following is an instance of (UI):

(3.1) $(\forall x Px \vDash Pa) \models$ (Everything is black $\vDash$ The Munttoren is black)

Besson (2019, forthcoming) claims that the above is not an application of (UI), but is rather an instantiation of a schema. Her argument relies on syntactic difference. While (UI) requires a universal statement as its premiss and then predication of the property of a single object, the above transition is purely syntactic. We simply replace the schematic letter P with the predicate *is black*, and the letter a with *The Munttoren*.

I agree with Besson that the above is a schematic instantiation, instead of an application of the (UI) rule. However, I believe this indicates a problem with the formalization of the argument, rather than with the (AP) itself. To appreciate that, consider again the rule we give Harry in natural language:

> *Universal statements imply each of their instances.*

The rule, I believe, is far from being schematic, at least the English version of it. Besson's conclusion seems to me to be the result of an over-reliance on formalisms. Therefore, Birman, or someone sympathetic to her argument, has an easy reply to Besson. To accommodate the schema reply, we can re-model the (AP) in order to completely remove the presence of schemas. To do this, we can follow the formalization of the (UI) rule due to Shaw (m.s. §5.1):

$$\forall x \forall y(\text{UIForm}(x, y) \wedge \text{True}(x) \rightarrow \text{True}(y))$$

Here, the predicate 'UIForm' applies to a pair of sentences in case one is a universal and the other is an instance of the universal one. The above formalization makes no use of schematic letters, unlike the (UI) rule. Both x and y are bound variables, bound by universal quantifiers. Therefore, in order to instantiate them, one will have to perform (UI), thus bringing back Birman's conclusion. Besson's argument from syntactic competence can be applied, but only to the conclusion that Harry will be able to apply 'UIForm', since recognizing that an argument is of the form required by (UI) can proceed in terms of pattern recognition.

But, one could reply, we are missing the point of Lewis Carroll's story, because the principle given above is just that: a principle, and not a rule. This reply, while fair, seems to nonetheless miss the point. Birman (2024, §3) discusses this line of thought:

> For someone who does not infer in conformity with the UI pattern, it is irrelevant whether it is given as a principle or rule. They both become a 'super premise'; one, in fact, that Harry won't know how or where to apply (p. 50).

Then, supposing we give Harry the rule-version of the (UI) statement:

> *From a universal statement, infer its instance*

The reasoning Harry will have to follow will take the rule as a further premiss (since, remember, he has no competence with (UI), so at first he will have to consciously think of the rule somehow). Both Birman (2024) and Devitt and Roberts (2024) agree on the fact that applying a rule requires recognition of some conditions of application, which very much resembles (UI). For some condition C, you must recognise that you are in an instance of C in order to apply the rule.

Of course, while the (UI)-principle stated before and the rule-version differ in phrasing, it can be nonetheless argued that they are more alike than one thought. In both cases, *some* competence with a conditional mode of thinking is required. For the principle, detaching the consequent requires (MP), while for the rule, you must recognise you are in some circumstance and then detach the imperative, which is basically no different than (MP). No accident that this is usually called the *Modus Ponens model of rules* (Wright 2007, 2018).

In the last chapter I will attack this conception in more detail, but for now I believe this is enough to do away with Besson's proposal. We can very well formalize the (AP) without using schemas, and while Harry will use some syntactic competence, this will be only in order to recognise arguments which have the form of (UI). Performing (UI) will still require some competence with (UI). As spoiler, the way I will attack this is by appealing to what is usually called *imitative learning* and goal inference.

TAKING STOCK

I have mentioned in passing what I think each solution gets right, and how I aim to improve on them. Here's an overview:

- From Devitt and Roberts (2024) I take the idea that Harry should undergo some form of training, but not one that amounts to conditioning
- From Chudnoff (forthcoming) I take the idea that some cognitive process should happen in order for Harry to truly be able to infer in accordance with (UI)
- From Finley (2024) I take the idea that some practice is necessary in order for Harry to understand the (UI) rule
- From Besson (2019) I take the idea that Harry should be able to engage some faculties related to pattern-recognition, but that on its own will not be enough in order to adopt (UI).

CHAPTER 4

Not understanding and lacking a concept

I claimed that Harry failed to adopt (UI) in virtue of not understanding the rule. Now it is time to substantiate my claim. The first half of this chapter will be spent motivating this diagnosis. Thus, I aim to show first why rejecting **Understand** is not an ad-hoc move to solve the (AP), but a principled, motivated rejection in its own right. The second half of this chapter will be focused on the problem of concept possession, arguing towards the conclusion that by raising the (AP) to the meaning/conceptual level we will be in a better position to help Harry, which will then be the topic of the next chapter. I will assume, and I will motivate this assumption in this chapter, that Harry lacks a proper understanding of the concept All. This will proceed by arguing that other possibilities about how Harry could have a partial understanding are unlikely. Therefore, we need to first help Harry acquire this concept, and this will be done by immersing him in a linguistic practice. What I will highlight is that it is impossible for someone both to have a proper grasp of a concept, while not being able[1] to use the respective concept. I believe that the intuition running through this chapter, while usually associated with inferentialism (Boghossian 2003), might be an attractive diagnosis in its own right. This is what this chapter aims to show.

[1] Where I take 'able' in the sense of 'having the ability', in order to block absurd conclusions such as someone not having a proper grasp of a concept because they are asleep.

4.1 Birman on inferentialism and concept learning

I will first walk through Padro's thoughts regarding the rejection of **Understand**. The rejection of **Understand** is, on Padro's (2015) account, a "natural reaction" (p. 158). Claiming that Harry did not apply the rule in virtue of not understanding it would not lead to the (AP), as you cannot adopt, in Padro's sense, something you do not understand. Similar situations arise when someone denies "Bachelors are unmarried men". We would claim that they simply have failed to understand what "bachelor" means. Then, the conclusion that Harry lacks the concept All is straightforward, as mentioned by Padro herself:

> Assuming the concept 'all' is constituted by its introduction and elimination rules, Harry's inability to infer in accordance with UI can be taken to support such claim. For, if the UI rule partially constitutes 'all,' not being able to follow it would indicate that he does not possess the concept (Padro 2015, p. 161).

What, then, prevents Padro from fully endorsing this solution? On the face of it, it assumes that concept acquisition consists of explicitly learning the rules for using a concept, which is highly unnatural. While the concept All can be seen as being constituted by its introduction and elimination rules, *learning* it is not as straightforward as learning just its rules. This way of learning concepts rightly fails in helping Harry. The view of learning concepts which I reject seems to be implicitly endorsed by Birman when she says that

> [w]e should be mindful, however, that claiming that Harry isn't in a position to accept the UI rule because he doesn't understand it, far from solving or dissolving the AP, just pushes it back to a reformulation in terms of meanings. Could we get Harry to 'adopt' the meaning of (or concept) 'all'? (Birman 2024, p. 54)

Thus, it seems that Birman is implicitly raising an Adoption Problem at the level of meanings/concepts. Let's formulate this new Adoption Problem explicitly, following the original statement of the (AP):

> Concept-(AP): certain basic concepts cannot be adopted because, if a subject already possesses the concept, no adoption is needed, and if the subject does not possess the concept, no adoption is possible.[2]

[2] If one is convinced by J. A. Fodor's (1975; 2008) arguments against concept-learning, then Concept-(AP) could be restated in terms of 'triggering' conditions (see J. A. Fodor 1981). Then, the problem is whether an explicit statement of a rule can cause a concept to become 'activated'. I will set this aside, and take on the assumption that concept learning is possible.

While Birman does not explicitly state the new problem,[3] it is clear from the context that she has an analogous version in mind when asking how Harry could adopt a concept. The problem with this analogy is, as mentioned, that she believes concept acquisition proceeds in terms of explicit rule-learning. Birman's main argument is to claim that prior acceptance of the (UI) rule is partially constitutive of the meaning of All, and thus if one cannot adopt the rule, how could one adopt the concept? This assumes that we acquire concepts by learning rules for using these concepts, which, if this were the case, would indeed spell trouble for my solution. An idea similar to Padro's is supported by Boghossian and Wright (2024, p. 99), who claim that any rule containing All can only guide Harry in accordance with his own understanding of the concept, and not a new one. However, as I will maintain, concept learning is inherently different from learning rules for logic, or indeed for concepts we already possess. Let me explicitly spell out her argument:

Concept-(AP) argument

(1) (UI) is constitutive for the concept All
(2) If one lacks (UI), then one lacks the concept All
(3) Adopting (UI) is necessary for acquiring the concept All
(4) Adopting (UI) is impossible
(5) ∴ If one lacks (UI), then one cannot acquire the concept All

I will deny (3), i.e., I will argue that one does not need to adopt a rule in order to acquire a concept, but that concept learning proceeds in a substantially different manner. My solution manages to preserve assumption (1) above, which is a key tenet of inferentialism,[4] while denying that learning a rule is how we *acquire* a concept. I will say more about this at the end of this chapter, and in the following one.

The connections between **Understand**, the (AP) and Concept-(AP) are thus as follows:

- Rejecting **Understand** means that Harry does not understand All (straightforward, and I will motivate this step in the rest of this section)
- Rejecting **Understand** raises Concept-(AP) (Birman's conclusion)
- If Concept-(AP) can be solved, then rejecting **Understand** solves the (AP) (straightforward).

[3] As a side-note, it will be highly likely that what it means for a concept to be basic might not be the same as what it means for a rule to be basic. For example, we can always formulate straightforward rules for complex concepts (e.g., rules for using the Sheffer stroke), and likewise complex rules for simple concepts (e.g., rules for the behaviour of numbers). However, I do not think much depends on a particular notion of 'basicness'. We can take the notion of basicness relativized to a certain community. A more detailed discussion is beyond the scope of this thesis.

[4] For a survey see Murzi and Steinberger (2017).

- Concept-(AP) can be solved (In the next sections I will lay down the plan for how I plan to approach this). In terms of the assumptions given in the first chapter, I will reject **Immediate** for Concept-(AP).

In the rest of this thesis I will argue that (i) we have a principled argument for rejecting **Understand** and (ii) that we can solve the Concept-(AP), thus avoiding replicating the (AP) at the concept/meaning level. My argument for (i) will be the subject of this chapter, and will proceed as follows, with each step of the argument having its own (sub)section:

(1) If someone does not assent to paradigmatic instances of the use of a concept, then the person does not understand the concept the way we do
(2) ∴ Harry does not understand the concept All the same way we do
(3) If someone does not have full understanding, they could have partial understanding of a concept: e.g., being able to make observational statements, or assenting to some of the constitutive rules of a concept
(4) I claim that this partial understanding does not suffice to individuate a certain concept. Purely observational statements without inferential capability do not count as concept possession, while partial assent underdetermines which concept is possessed.
(5) ∴ Harry lacks the concept All.

4.2 When can we say that someone does not understand?

The idea that Harry does not understand the concept All due to not using it properly seems an attractive diagnosis. However, is it simply an *ad hoc* move to avoid the (AP)? I claim it is not. In this section I will give a principled response based on prevalent views from the literature, the idea that someone who does not use[5] canonical instances of an expression does not understand the expression in question. Thus, charging Harry with not understanding All or lacking the concept is not only an intuitively good answer, but the right one. I will then discuss a well-known counterexample to this view, Williamson's (2022) criticism of Understanding-Assent Links, and argue that we can overcome his critique. My argument extends those by Incurvati and Schlöder (2023, §2.5), Warren (2021) and Wikforss (2009).

It is notoriously difficult to spell out what understanding an expression consists in (for a survey, see Hannon 2021). However, it seems easier to spell out what misunderstanding would looks like. In other words, we are looking for a minimal principle which, if violated, would safely charge someone with misunderstanding the expression used.

[5] In this section, I will be somewhat sloppy between using, assenting, recognizing an instance, etc., since I believe that whichever one we would ultimately prefer in formulating the principles at hand, Harry is capable of neither.

This issue is more tractable, and fortunately, there is one such principle already endorsed under various forms in the literature (Brandom 1993; Frápolli 2023; Hart 1961; Incurvati 2020; Leech 2015; Williamson 2000). For someone to understand an expression, they are required to accept/recognise/be able to use some easy instances of this expression:

> being able to perform certain "easy" inferences is constitutive of understanding a given logical constant and therefore is necessary for believing *any* sentence containing that constant (Cherniak 1986, pp. 39-40).

> one cannot regard anything which is recognizably an instance of the relevant inference pattern as unsound without convicting oneself of misunderstanding. (Hale 2002, p. 290)

> [C]onsider the case in which the expert's inferential behaviour shows that they do not accept common-or-garden applications of *modus ponens* [...] In this case, we would be entitled to question their grasp of the meaning of *if*, their expertise notwithstanding (Incurvati and Schlöder 2023, p. 58)

The common thread in the above quotes is that understanding implies at least recognizing/assenting to some "easy" instances of the expression at hand. We do not need full overlap in meaning in order to use the same expression, but some minimal overlap is needed in order to not talk past each other (for a a defence of this idea see Drobňák 2024). Let me phrase this principle in a more straightforward manner:

> **Minimal-Understanding:** Someone does not understand an expression if they do not accept/recognise/use[6] canonical instances of the expression.[7]

To give an example, someone who would claim that two sets X and Y have the same members but are not the same misapplies the concept Set since extensionality is a core property of sets (Incurvati 2020, p. 11). Thus, it seems that, for some class of "easy" or "recognizable" instances of a concept

[6] I take this acceptance in dispositional terms, meaning that we do not consciously accept all of these instances, but that when presented with them, we're disposed to accept them. Importantly, I am not claiming that they must accept the canonical instances as *truth-preserving*.

[7] I use the neutral term 'expression' to try and capture any sort of thing about which we could plausibly talk about someone's understanding of the respective thing. With this I aim to include semantically-relevant expressions such nouns, adjectives, rules, etc., and exclude morphemes.

or an inference pattern, not accepting some instance[8] from this class charges one with misunderstanding, whereas for harder instances one could deny instances while still plausibly retaining understanding of the concept involved. The condition is minimal, meaning that I do not believe endorsing canonical instances is sufficient for understanding. Let me give some examples (and, as will become clear in the rest of this section, examples have an important role to play in Harry's case). Let's take the concept Bald. Everyone would agree that Michael Jordan is bald and that Mick Jagger is not. However, how would we react if someone were to call Mick Jagger 'bald'? Supposing they are not using the word figuratively and they can perfectly see Mick Jagger's hairline, then it is likely we are going to say that they are simply misunderstanding what the word means (compare with Grice and Strawson 1956, pp. 150-2).

Before we move on, a note on terminology. I have chosen the term *canonical* as opposed to *easy* for one important reason. There are many 'easy' inferences we could draw using an expression, but not all of them are important for the meaning of that expression. For example, a chair can be made of wood, plastic, metal, or can have a certain shape. However, these properties do not seem to be canonical for something to count as being a chair. On the other hand, we would find it odd if someone would not understand that a chair is for sitting. As Rabin puts it, "If one fails to realize that chairs are for sitting, one fails to master Chair, no matter how many other Chair-truths one knows. Mastery is tied to grasping key insights about the concept" (Rabin 2020, p. 630). While, for logical concepts at least, I think there might be considerable overlap between what we consider 'canonical' and what we consider 'easy', I opt for the former terminology since it is more readily applicable beyond logic.

Let me also give an example from logic. Let us consider the concept All, the one Harry has so many issues with. As noted in Chapter 2, the concept appears in multiple expressions, such as *everything*, *everyone*, etc. However, one could understand the concept without also understanding all the expressions in which it figures. Indeed, this is why we are, in Harry's case, focusing on the (UI) rule. Similarly to how a chair is more tightly linked with its sitting purpose than with its shape, so is the (UI) rule more intimately linked with the meaning of the concept All than other expressions (see Dummett 1991, p. 275; also Edwards 1995; Goodman 2024, p. 88). Let's take the following as one canonical instance of (UI):

$$\text{Canonical-(UI)}\ \frac{\text{Everything exists}}{\text{I exist}}$$

[8] Note that I do not claim that endorsing these instances is how we settle on the meaning of some concept. I am claiming something much weaker, namely that no matter how we have settled the meanings, such-and-such instances seem highly acceptable in our current linguistic community. This is in line with suggestions made by Wikforss (2009) in that the understanding should be construed holistically.

By contrast, the following seems to be a rule in which All features as an essential component, but which we do not so readily call canonical for the meaning of the quantifier (Brîncuş 2024):

$$\omega\text{-rule}\ \frac{Pa \qquad Pb \qquad \ldots}{\forall x Px}$$

Importantly for my proposal, in the case of the ω-rule, someone can be deemed competent in using the concept All both if they deny or if they accept the ω-rule. On the other hand, someone who denies that Canonical-(UI) is a canonical use of the concept All might be charged with misunderstanding. Modifying a claim made by Leech regarding the conditional, we could say that

> if a logical connective doesn't behave so as to validate at least simple instances of [(UI)], then it doesn't count as a [universal quantifier]. [...] [The language] contains a logical constant or connective with various properties, but that constant should not be counted as a [universal quantifier] (Leech 2015, p. 25).

Thus, one can either endorse or reject ω-rule without, so long as they accept Canonical-(UI) or some version of it, being accused of not understanding the concepts involved. Before delving into Williamson's critique, I want to pre-emptively discuss some possible worries related to my proposal.

One can complain that what makes something a canonical instance depends on one's linguistic community. If some community would, for example, find Fermat's Last Theorem immediately obvious as the Mathoids in Berry (2013) do, then such a community might deem the one-step inference from Peano Axioms to Fermat's Last Theorem as a canonical inference in the case of arithmetic, which is not the case for us. This spells problem if we think of canonical inferences as being valid inferences and thinking that they have a special status when it comes to *a priori* justification. While this indeed is a thorny point for those that want claim that such canonical instances have a special epistemological status (for example, Enoch and Schechter 2008; Wright 2004b), my proposal does not mention validity at any point.[9] All I claim is that someone who fails to recognise a canonical instance can be charged with misunderstanding the expression as currently used. The Mathoids might capture the same facts as we do, but possibly using different concepts. They can also charge us with misunderstanding arithmetical expressions due to not assenting to their canonical inference to Fermat's Last Theorem, but this misunderstanding would be relative to *their own* meanings and concepts. Thus, different communities can take radically different instances as canonical, and thus understanding should be relativized to what a certain community means

[9] Indeed, Warren (2021) also draws a distinction between understanding an expression, and the assumption that this understanding does epistemological work.

in using a certain expression. This is not a problem for the (AP), because it seems that Harry is part of no community which uses All, and what we are trying to do is to get him to adopt *our* understanding of All. Thus, if so desired, we can restrict understanding and canonical instances in the way I've mentioned, i.e., relativized to a certain community.

Another worry could be that, despite assenting to different canonical instances, two speakers can nonetheless have the same understanding of an expression.[10] This is an example due to Cherniak (1986), where someone who takes other easy instances as basic might nonetheless mean the same thing as us: "[a] creature that cannot [infer according to disjunction introduction and disjunctive syllogism] may have enough *other* inference ability so that we can distinguish 'or' from 'and' for this creature" (Cherniak 1986, p. 41). I think we can accommodate this case as follows. Suppose we have a translation[11] in place between us and the creature. Perhaps it cannot assent to our rules as they are, but it could assent to the way we might translate the rules in its language. If it accepts the translation of these canonical instances, then the divergence is only apparent, and we end up meaning the same thing. However, if disagreement occurs on such a canonical instance, then we would have reason to doubt that the creature actually uses the logical constant as we do. For example, if someone says *La neige est blanc*, I could not assent because I don't speak French. However, if someone translated it for me and then read it as *Snow is white*, I would assent, and thus there would be no divergence in meaning. The lesson of this paragraph is that agreement on canonical instance is not needed from the get-go, i.e., that both speakers take the same canonical instances as canonical, but that they do not disagree on each other's canonical instances. This case might allow for initial apparent disagreement, but later on both speakers can realize that they have been speaking about the same thing, only using different words (for a similar idea in terms of verbal disputes and Stalnakerian common ground see Miller 2022). To invoke the example of the Mathoids, if they would take the inference to Fermat's Last Theorem as valid, and if in our community the Theorem would have actually been proven false, *then* we could raise the issue that there is a difference in concepts between us and the Mathoids.

Lastly, one could wonder if my view is simply a watered down version of what is usually called the *Prototype Theory of Concepts* (see Laurence and Margolis 1999), the view that representations usually include typical properties

[10] Or, to use the terminology of Evnine (2008, p. 26), they can possess different but logically equivalent concepts.

[11] A question can be raised about what sort of conditions we place on such a translation. If liberal enough, we could in principle always find a translation that maximises agreement despite there being substantial disagreements between speakers (e.g., translating everything someone says as '2+2 = 4'. Perhaps a criterion such as definitional equivalence could be introduced, and indeed the way I present the scenario seems to point in this direction. However, a proper discussion of translation criteria is beyond the scope of this thesis.

of their elements,[12] such as someone being a Grandma typically being an old woman, with gray hair, sweet and walking with a cane. However, a grandma can also be someone much younger, and without a cane, etc., in other words someone untypical can also fall under a certain concept. The analogy with my canonical statements should be clear, with these prototypes standing is as my canonical statements. I think the similarity, while still there, should not be a concern. I spelled out my proposal in terms of assent, and not in terms of actual mental representations. In other words, it is a matter of actual agreement whether we mean the same thing, and not a matter of mental representation. If we have enough common agreement, then we are not talking past each other. To keep the grandma example, so long as we agree that the old lady is a grandma, we might have enough common ground to walk on. In terms of conditions, it seems that having the same prototype might guarantee sameness of meaning, but that assenting to canonical instances might lead to different prototypes, and thus different representations. With these preliminary discussions behind us, I now turn to a more significant concern, namely Williamson's (2022) critique of the so-called understanding-assent links:

> (UAL): Necessarily, whoever understands an expression[13] p assents[14] to it.

The (UAL)s are what put forward the seeming obviousness of the examples discussed in the previous paragraphs. For some expressions, one must assent to them if one understands their meaning, and not assenting would charge one with not understanding the expression. If the (UI) rule is, as Birman claims, constitutive for understanding All, then an inferentialist seems committed to the following

> (UA-All) Necessarily, whoever understands (UI) assents to it.

However, counterexamples seem easy to find. The most common example is that of free logicians, who deny that (UI) applies for non-denoting terms (Nolt 2021), and as such reject the (UI) rule in full generality. Similarly, Gauker (1997) rejects the validity of (UI) by employing a notion of context. Then, not only is it possible for someone to understand All without assenting to (UI), but there are actual examples of such rejection. Can we hold on to

[12] The discussion is not limited to cognitive science, see for example Rumfitt (2015, Ch. 8) who uses the idea of 'poles', which are pretty much typical examples, in order to talk about vague terms.

[13] I will go in tandem with both versions of understanding-assent links, those for expressions and those for thoughts (see Williamson 2022, pp. 75-6). My discussion will bear on both, but certain aspects are easier to emphasise in certain renderings.

[14] We can ignore the supposed epistemological work (UA) is supposed to do in justifying *a priori* knowledge. Harry's case is simply a case of recognizing some instances that we take to be representative of a concept, but we can be silent on whether these instances are indeed valid.

(UAL) and claim that these experts don't actually understand what they're talking about? Williamson chooses to block (UAL) instead. In the rest of this section, I will discuss in more detail Williamson's argument and show how it might be overcome based on my endorsement of canonical/non-canonical instances. In other words, we can both keep (UAL), and not charge the aforementioned experts with misunderstanding. I will weaken (UA-All) in the course of my argument, but the end result will still be enough to offer a rebuttal to Williamson.

We can render Williamson's argument more explicitly, following Warren (2021, p. 9225), as follows:

(1) It is possible that some *S* fails to assent to *p*
(2) *S* means by *p* the same we mean by *p*
(3) ∴ (UAL) does not hold.

On Williamson's account, someone like McGee (1985) is a counterexample to (UAL) due to his rejection of (MP). Williamson's challenge seems to threaten my principled conclusion that Harry does not understand All, since denying (UAL) leads to an in-principle rejection of **Minimal-Understanding**, since it is possible to think of someone who, for various theoretical reasons, might reject garden-varieties of (UI), even Canonical-(UI). This leaves open the possibility of Harry understanding the expression without using or accepting it.

In light of Williamson's challenge, three approaches seem available. One can either (i) accept Williamson's point, but deny that it applies to Harry's case, (ii) refine the notion of **Minimal-Understanding** or of (UA) such as to circumvent Williamson's critiques or (iii) challenge Williamson's argument, and defend the possibility of (UAL) (this is the path taken by Warren 2021). I will first discuss (i) and (ii) which I will reject, before endorsing (iii).

(i): *Accept Williamson's point, but deny that it applies to Harry's case.* This solution seems to be based on a very weak notion of assent. In the scenario, Harry accepts the (UI) rule as a good rule, and this is indeed a necessary step in Padro's recipe for adoption. Thus, it seems that he can assent without needing to understand the rule, and this renders the (UAL) link trivial in this case. We can assent to things we do not understand, such as when I accept a mathematical theorem my professor wrote on the board. However, we can refine assent to mean *assent in virtue of understanding* the expression (Warren 2021, p. 9222 makes a similar comment), in which case this solution fails. This solution is thus based on a shallow understanding of assent, and we can reject it.

(ii): *Refine the notion of* **Minimal-Understanding** *or of (UA) such as to circumvent Williamson's critiques:* Changing the phrasing of (UA) seems to beg the question against Williamson. The most straightforward way we could try to revise (UA) is by explicitly embedding the clause that the expression must be canonical:

> Canonical-(UA) Necessarily, whoever understands a canonical expression p assents to p.

However, Williamson's critique still applies in this case. Remember that his argument is based on the idea that for any given expression p we can conceive an agent who fails to assent to p, and there are no limits on what kind of expression p might be. Consequently, p can also be canonical, and Williamson's argument will still hold. Thus, it seems that restricting the (UA) clause to a certain class of expressions fails to evade the critique and Williamson's argument applies to all expressions.

Can we revise **Minimal-Understanding** instead? I believe the claim, by already being minimal, captures something that seems necessary for understanding an expression. As in the previous paragraph, we can try to restrict further what is required for understanding (e.g., only one canonical instance), but Williamson's argument will continue to apply for virtually the same reasons as before. As long as we require assent to *some* sentence, Williamson's argument will still have force. With this in mind, I now turn to what I find a more plausible approach, namely:

(iii) *Challenge Williamson's argument:* The argument has two premisses which we could dispute, either the possibility of some S failing to assent to p, or that S would mean the same thing by p as us. I take the first premiss to be in order. Indeed, we could think of any number of far-fetched theoretical[15] reasons someone might fail to assent to some p. Thus, to reject Williamson's argument we need to reject the second premiss, i.e., that someone who does not assent to some p has the same meaning of p as us. The line of reasoning I will be using is one in term of pre-theoretical commitments and the issue of 'talking-past-each other' (Brun 2022; Carnap 1950). What I want to argue is that we should go one level higher, namely to the level of concepts. I believe there is the right place to make the case for (UAL). Let's write this down more explicitly:

> (UAL)-Concept: Necessarily, whoever understands a concept C assents to some canonical use of C

The move to concepts is legitimate on Williamson's terms, and is agnostic as to how fine-grained we individuate concepts (see Williamson 2022, pp. 118-122 for a discussion). Naturally, I do not think that all concepts will exhibit such an (UAL),[16] but we do not even need to push that far. I believe that someone who understands a concept will assent to more than a single

[15] I say theoretical to avoid accidental lapses in understanding, such as when someone fails to endorse (UI) because they don't know what 'exists' means. I could have said principled reasons as well, the main idea is that the reason for rejection has entirely to do with the expression itself, and not with considerations about other elements.

[16] For example, for concepts such as Art, see Gallie (1956).

canonical use, but let's concede as much ground as possible to Williamson.[17] In line with my discussion of the previous pages, I also take it that some community might choose a particular set of canonical-instances, different from another community, but the pattern should be there, i.e., assent to some canonical uses of C. To get an idea as to how I aim to defend (UAL)-Concept, recall Rabin (2020) saying that someone who does not realise that chairs are for sitting fails to understand the concept Chair. In Williamson's example, any expert could be brought forward who fails to assent for whatever far-fetched theoretical reasons to some expression *Chair X is for sitting*, while still plausibly retaining understanding of the concept Chair. However, if we bring forth someone for whom no conceivable chair is fit for sitting, then that sort of deviance will break the seeming of understanding.[18] That deviance is too extreme for understanding to be retained.

Let's start with a simple example. It is clear (at the moment of writing) that Sweden is a democracy and that North Korea is not a democracy. Thus, the way we use the concept of Democracy includes commitment to labelling Sweden a democracy and to not labelling North Korea a democracy. Now, suppose some *S* claims *North Korea is a democracy*, and we translate homophonically, since they "intend their words to be understood as words of our common language, with their standard English senses" (Williamson 2022, p. 91). Assuming they are aware of the other relevant facts, e.g., about North Korea's politics, its history, and so on, it then seems plausible that *S* uses the concept Democracy in a different way than us. The examples of Sweden and North Korea are what I've called so far canonical instances, in this case Sweden being an example and North Korea a non-example.

Supposing *S* has access to the same facts[19] as we do (e.g., we are both read an objective political report on both countries), the violation of such canonical commitments could be best explained by claiming misunderstanding of the concepts involved.[20] They simply do not mean the same thing as we do. However we might want to precisify the concept Democracy or spell out necessary and sufficient conditions, due to our current commitments any plausible and acceptable effort must declare Sweden a democracy and North Korea not a democracy. Note that I do not claim that Sweden must be a "democracy" in an absolute sense. It would be possible that we might change

[17] This concession goes strongly against Quine's *change of logic, change of meaning* thesis (Quine 1986; Warren 2018), but I want to show how we can overcome Williamson even without endorsing Quine's very strong change of meaning thesis.

[18] My line of argumentation is slightly inspired by Putnam (1978).

[19] Important to note that I do not consider facts as involving the concept at issue, in a move similar to Chalmers and Jackson (2001, p. 323), cf. also M. Balcerak Jackson and B. Balcerak Jackson (2012, p. 196). Then, competent concept use requires both concept possession and sufficient information about the world. I maintain the world-knowledge fixed for all subjects in order to zero in on conceptual issues.

[20] Similarly to the 'orangutan' in Burge (1979), the 'hippopotamus' in Davidson (1968), or the 'three-year old adult' in Grice and Strawson (1956).

the meaning of Democracy in a surprising but acceptable way, and have Sweden not qualify as a democracy. However, such cases seem to be more radical changes of meaning which happen at the level of the community and which would lead us to changing the concept itself. Thus, it would be more accurate to say that whatever it is we mean by Democracy *now*,[21] Sweden is a democracy and North Korea is not, and that denying these basic commitments means that one does not understand the way we use the concept.

As mentioned before, according to **Minimal-Understanding** principled rejections in non-canonical instances, such as McGee (1985) on (MP), are allowed, and experts will still mean the same thing by Implication as we do. Rephrasing the democracy example in terms of implication, we can say that the move

$$\frac{\text{If it is raining, then the streets are wet} \qquad \text{It is raining}}{\text{The streets are wet}}$$

is a canonical example of (MP), whereas one of the examples from McGee (1985, p. 462)

$$\frac{\text{If a Republican wins, then if it's not Reagan it will be Anderson} \qquad \text{A Republican wins}}{\text{If it's not Reagan who wins it will be Anderson}}$$

is, while still formally (MP), not such a canonical instance. Then, someone who denies the first and accepts the latter might very well be said to not understand how to use implication. Of course, I do not mean that someone who has a momentary lapse in reason and does as described at first, but then corrects their views, lacks understanding. They do understand, and the correction is due to such understanding. I mean someone who regularly, consciously, and honestly denies the first and accepts the latter can very well be said to not mean the same thing as we do when using the implication.

Williamson's point in arguing against (UAL) is that enforcing (UAL) precludes disagreements, and he has in mind disagreements over conceptual truths. Thus, language must not impose conditions of understanding as that could limit the possibility to rationally doubt certain expressions. However, I believe that Williamson's recipe for acceptable deviance fails because it does not consider extreme enough deviances. What I mean by that is that the cases considered by Williamson are indeed intelligible and intuitively right (e.g., someone denying that *Every vixen is a female fox* due to a peculiar view of quantification), but that simply means we have to push further. This is what I want to do now.

Here's one point I am willing to concede to Williamson for the sake of argument. It could very well be that for any one expression we can find

[21] To give another example, the concept of Democracy as it was understood in Ancient Greece would not apply to most Western countries, since Greek democracy was not representative, but direct.

someone who will disagree, and yet have enough agreement in other parts of language for them to still mean the same thing as us. One single deviance is, indeed, not enough, but a concept-relevant pattern is needed, indeed one that goes against the whole practice of a community. In other words, I agree that "sufficiently fluent engagement in the [linguistic] practice can take many forms" (Williamson 2022, p. 128), but I disagree that there will be "no single core of agreement" (ibid.). The issue is that this core will not be a single sentence or thought, but it must be a concept-relevant pattern of use. Here I wholeheartedly endorse Kevin Scharp, who says:

> [S]omeone who possesses a given concept (is competent with a given word) accepts some of its constitutive propositions (sentences), but need not accept any particular one of its constitutive propositions (sentences) (Scharp 2013, p. 55).

One concept for which I aim to exemplify (UAL)-Concept is the concept Truth. I believe a similar argument will then follow by analogy for the concept All, but the current literature is more telling for my argument in the case of Truth.[22] Then, most people will agree with at least some instances of the introduction and elimination rules for Truth:

$$\text{T-In}\ \frac{\varphi}{\text{True}(\ulcorner\varphi\urcorner)} \qquad \frac{\text{True}(\ulcorner\varphi\urcorner)}{\varphi}\ \text{T-Out}$$

This seems intuitive enough, and most theories of truth validate at least *some* of these rules. Indeed, most go as far as to say that for something to count as our Truth, it *must* conform to these rules (Scharp 2013; Tarski 1943). Of course, some do restrict the applicability of these rules in order to avoid Liar-paradoxes (for a meta-analysis, see Horsten and Halbach 2015). These rules seem to codify something central to our use of the concept. Nonetheless, even here Williamson has his deviant champions. In the case of Truth they are known as *alethic nihilists*, consisting mainly of the pair Gamester (2023) and Liggins (2019). I believe that their case is enough a deviance to claim a difference in meaning. To show that, I will make a point from the literature on what is usually called *merely verbal disputes* (Chalmers 2011; C. S. I. Jenkins 2014; Miller 2022), following an argument due to Asay (2021).

The nihilist denies that there are any truths, rejecting "all instances of (T-In)" (Gamester 2023, p. 318) for principled reasons, namely to avoid Liar paradoxes without arbitrarily restricting the rules for using Truth. In other words, no matter what I replace for φ in T-In, the nihilist will deny it. Thus, the nihilist is a textbook example of one of Williamson's deviants, who has principled reasons for rejecting something we might deem intuitively

[22] Since I do not believe there is any existing position which rejects (UI) across the board, unlike the case I am about to present.

obvious. However, this is, I maintain, one deviance which goes too far. Due to the centrality of T-In to our understanding of Truth, someone who rejects *all* instances of T-In is too far gone to have the same meaning as us. For Williamson, the deviant must still speak the same language, with pragmatics allowing two speakers to still be mutually intelligible. Thus, someone who rejects T-In for the Liar sentence and someone who endorses it can still understand each other, and have a meaningful debate about how Truth should be understood. This is because "defenders of rival theories need not talk past each other as long as they share a suitable point of reference" (Brun 2022, p. 21). One can believe that 'Truth' is whatever the Pope is wearing. They could still manage to be intelligible so long as they assent to canonical instances of use.

Nonetheless, are there other commitments that the nihilist might have in order for us to still understand them? They are, for example, still committed to T-Out, so there is still some common ground. Might this suffice? I do not think so. While T-Out is also something we consider canonical for Truth, it is not sufficient to be canonical on its own. Our commitments to the use of the concept make it so that T-In must also be endorsed. Thus, it is the pair of these rules which allow us to distinguish the concept Truth from another concept which might only conform to one of the rules (for example, Scharp's notion of descending truth which only conforms to T-Out, see Scharp 2013, Ch. 6). I will say more on this in §4.3.2 below, but this should give some reason to think that due to the way we use Truth, assent to only T-Out will not be enough to interpret the deviant as talking about Truth. This is a situation which fits Warren's point that "if deviance is extreme enough, the push toward homophonic translation is weakened to the point of non-existence" (Warren 2021, p. 9228).

Putting the point in terms of translation as Warren does gets to the crux of the matter. These constitutive principles are determined by our use of the concept. If we were to have used the concept otherwise, other principles would have replaced the current ones. Then, these constitutive principles are not only used for individuating the concept and carrying out conceptual analysis, but importantly we use these constitutive principles as part of our conversational record: "I think it is an empirical fact about our linguistic practice that we use constitutive principles to guide interpretation" (Scharp 2013, p. 38). Someone, like the nihilist, who denies T-In across the board will not be interpreted as meaning Truth by 'Truth' by any competent member of the community. Their insistence to the contrary only makes matters worse.

What I mean by that is that the nihilist, by doubling down on their conviction that they are still talking about the same thing, are in fact halting prospects for advancing philosophical discussions. Supposing we are interested in the nature of Truth, as any philosopher would be, what do we gain from the nihilist's insight? I believe that taking the nihilist seriously would make the ensuing discussion a merely verbal one. When the non-nihilist

says *Something is true* and the nihilist counters with *Nothing is true*, their disagreement will no longer be one over Truth, but one about 'Truth', i.e. about the use of the word. I am not saying that a meta-linguistic dispute is not something worth having, but I am saying that it is not worth having it while under the impression that it is a factual dispute. In our community, the bond between Truth and T-In is strong enough, both in-use and upon reflection, that any point to the contrary is changing the subject.

Since Williamson's main reason for rejecting (UAL) is as to not enforce restrictions to the scope of disagreement, it seems that allowing the nihilist to count as still using the concept with its usual meaning will, unfortunately for Williamson, be one such disagreement that will be barred. If "[u]nderstanding words in a natural language has much to do with the ability to use them in ways that facilitate smooth and fruitful interactions with other members of the community" (Williamson 2022, p. 99), then the aforementioned case seems anything but a fruitful interaction. If agents disagree on canonical instances but nonetheless have the same meaning and access to the same facts, any kind of meaningful disagreement seems impossible, and they would be talking past each other.[23] To prevent this worry, it seems that agreement on canonical instances is needed:

> Less could suffice to ensure that debaters do not talk past each other, for example an agreement on paradigmatic instances and non-instances (Brun 2022, p. 6).

Am I not begging the question against Williamson by saying that the nihilist changes the subject? Not quite. As Williamson would have it, words get their public meaning from the linguistic community, and there must be some reason the community uses a concept. For instance, the function of truth is at least its use as a disquotational device (Picollo and Schindler 2018; Quine 1986). Then, the community needed a concept X which allows it to go from φ to X(φ) and back, and it just so happened that it named it 'Truth'. The community's need for such a disquotational device will not be superseded by claiming that we should stop going from φ to X(φ), but will simply attribute this function to something else, say Y. If Truth is the concept which is supposed to satisfy such a function, then the nihilist's argument will not dislodge this. The nihilist could argue that we should give up disquotation as a useful device,[24] and thus remove the need for it, but this is not what they are aiming at.

[23] The above worries can also be connected to recent issues related to change-of-subject in the conceptual engineering literature, see, for example, Burgess et al. (2020) and Cappelen (2018). Due to reasons of space, I will not pursue this line of thought in this thesis, but I believe it to be an interesting connection worth pursuing in future work.

[24] Indeed, this is how so-called 'ameliorative' projects proceed, with Scharp (2013) being a prime example.

I believe this is enough to establish that anyone who understands Truth must assent to some instances of T-In and T-Out. Now, I believe that a similar argument can be carried out for the concept All. Our use of the concept All makes it so that (UI) is constitutive for the concept, and perhaps this is due to the need of a community to be able to make a move from universal statements to particular instances. By repeating the argument of the previous pages, someone who denies (UI) across the board can be taken to change the subject, and thus not understand the concept the same way we do. This is why in order for someone to understand All they must assent to *some* use of (UI). Naturally, as mentioned before, I do not think this line of reasoning can be carried out just for any concept, but the ones

What was the purpose of the aforementioned argument? It seems intuitively clear that there is something wrong with Harry, but why should we go as far as endorsing canonical instances and rejecting Williamson's argument? As mentioned in the beginning, I intend my rejection of **Understand** to be a principled approach to overcome the (AP). In other words, I want my diagnosis to be solid, independent of other proposals. To that end, the aforementioned discussion regarding understanding seemed necessary in order to reach the principled conclusion that Harry, or indeed anyone that does not recognise easy instances, lacks understanding of the target concept. The conclusion would be weaker without endorsing such a connection, since then Harry could very well understand the rule despite not being able to perform inferences, just as Williamson's possible experts can understand without assenting. What I have aimed to show is that understanding a certain expression allows one to recognise canonical instances in which the target expression is being used. Lacking the ability to recognise such instances charges one with either lacking, or not using the concept in the same way we do, with both cases entailing misunderstanding.[25]

As one final point before I move on, as a final push against Williamson, let's take the debate on fundamentality in metaphysics as a more fleshed out case study in meaning-difference. Many metaphysicians endorse the idea that ordinary objects do not exist, and only, e.g., fundamental particles exists (Sider 2011). In other words, there are no tables, but atoms arranged table-wise. Then, is this a case of understanding, e.g., *This table exists* while not assenting to it, contrary to ordinary speakers? The case here seems to be immediately settled against Williamson, as these theories are best phrased in a theoretical language, and not in English:

[25]Importantly, my proposal is consistent with cases of misunderstanding such as Burge's (1979) patient who misunderstands Arthritis. Burge aims to show that one can possess a concept without having full mastery of its use, which is something compatible with my proposal, e.g., one can have the concept of Implication without being aware of McGee's counterexample. Likewise, one can have the concept Implication, but deem Affirming the consequent to be valid. In that case, we can re-apply the worries of the previous pages and doubt that the person understands the concept the same way we do.

> Perhaps my book, and other works of ontology, should not be interpreted as English, but rather as "Ontologese", a language distinctive to fundamental ontology, in which the quantifiers are stipulated to *mean something new* (Sider 2004, p. 680, my emphasis).

Then, there is no actual disagreement *in English*. The apparent disagreement is due to translating homophonically from this fundamental language into English, taking the fundamental '∃' to mean the same as the English *exists*. This is the same lesson as before, where we would be mistaken in translating someone's *if-then* as our *if-then* if they reject a canonical instance, or translating someone's use of *democracy* as our use of *democracy* if they claim that *North Korea is a democracy*. If Ted Sider would come in my room and claim that my bed does not exist, I would know he's speaking Ontologese, and I would know to not translate homophonically, but the same might not be true in the case of my non-philosophically trained friend. They would just claim he's blind or mad, in virtue of translating homophonically. I do not see this as any different to the person claiming that *Every vixen is a female fox* is false in virtue of a peculiar view of quantification. They mean something different by *every*.[26] Likewise someone who thinks a vixen is a male fox[27] clearly should be translated as saying *Every reynard is a female fox*.[28]

Let's recap. In this section, I have argued that rejecting **Understand** is not an *ad hoc* solution to the (AP), but a principled diagnosis in its own right. The way I argued for this was by considering what I called *canonical* instances, instances that can be thought of being so central to our understanding of a concept or a rule that denying them can charge one with misunderstanding. The last part of this section was focused on defending my proposal against Williamson's challenge to (UAL), and I have shown that his argument fails because a speaker with deviant enough patterns of use related to the concept at hand will fail to have the same meaning. Denying this conclusion would require Williamson to give up one of his own methodological goals, namely the possibility of disagreement. Should Williamson abandon the motivation from disagreement in order to accommodate sameness of meaning, then the rationale for denying (UAL)s also disappears. In any case, both of Williamson's possible replies seem unnatractive. In the following section, I will further refine my conclusion, considering ways in which Harry could have a partial understanding of the concept, and show that such partial understanding is insufficient.

[26] Thomasson (2014, pp. 24off) puts forward a similar line of thinking.

[27] I do think there is a difference between these two cases, with this latter one perhaps being more clearly a counterexample to Williamson. For more on this class of stipulative truths, see Soysal (2020).

[28] Where *reynard* is the posh term for a male fox, just as *vixen* is for female fox.

4.3 HARRY'S CONCEPTUAL RESOURCES

In the previous section, I claimed that Harry does not understand the concept ALL properly, where by properly I mean at least approximately the way we understand and use the concept by endorsing or recognizing what I called canonical instances of use. This section will discuss two ways we can make this lack of understanding more precise: either Harry has only a *partial* understanding of the concept, meaning that he might assent and use some of the rules we do, but not all of them, or that Harry lacks the concept completely. I will discuss both interpretations, and focus on the latter one, as I believe the way we learn concepts in such cases can also be applied in cases of incomplete understanding.

Harry does not understand ALL the same way we do, meaning that he does not use the (UI) rule. As mentioned in the previous paragraph, the possibilities are:

(1) Harry has a partial understanding of *all* as we do, meaning he assents, e.g., to the introduction rule for ALL: $Px \vDash \forall x Px$, or is able to correctly make true statements including the concept, such as *All mammals are vertebrate.*

(2) Harry has no understanding of ALL: he lacks the concept.

As a clarificatory point, when I mention *concept possession*, I am excluding mere *nominal* possession in the sense of Bealer (1998), such as when a child first hears the word *plus* and makes some mental space for this word, only associating it with the one sentence, e.g., all they know about ADDITION is that '2 + 2 = 4'. Such concept possession would not allow the child to perform addition, which seems to license the conclusion that the child lacks a proper grasp of ADDITION. I use *concept possession* more in the sense of Crimmins, who says that:

> 'Having the concept', I suggest, is simply having a *normal* idea; it is having an idea involved in normal beliefs and associated with normal recognitional abilities. Knowing about the potent solvent 'H_2O' is not sufficient for having the concept of water. One would have to believe that water is a clear, more or less tasteless, colorless liquid at room temperature, that it is safe to drink, and so on. One needs *normal* beliefs and *normal* recognitional abilities to count as having the concept (Crimmins 1989, p. 287).

I do not think much hangs in the balance when choosing one terminology over the other, but I state my use explicitly to avoid confusion. Further, to avoid this potential worry along the way, I do not claim that *full* understanding is required to possess the right concept (for a similar point see Siebel 2005). Then only logicians could possess logical concepts, which seems plausibly false. As mentioned, assent to canonical instances is a *prima facie* plausible

constraint to avoid talking past each other, but beyond that I do not impose other restrictions. Thus, when I mention that Harry lacks the concept, I mean that he lacks this minimal grasp of the concept. Going forward, whenever I say *concept possession*, I mean adequate grasp of the concept, and whenever I say *lacking a concept*, I include cases of inadequate grasp as lacking the concept. The (AP) could then be seen as a problem to the idea that one can transition from an incomplete grasp to an adequate one by being told a rule that would fill this gap, namely (UI). However, my rejection of **Understand** still stands, and even in this way of viewing the (AP), we can still raise the Concept-(AP) and solve it as advertised.

Assent to canonical instances is, as mentioned, necessary, but not sufficient for understanding a certain concept. If I hear 1000 doctors unanimously and easily diagnose a patient as having small pox, then I will assent to *Patient X has small pox*, and thus have what I have called a nominal possession of Small Pox. In this case, I would assent to a canonical instance without understanding the concept, and that is acceptable. What would count me as *not* understanding the concept would be denying *Patient X has small pox*, case in which I will, unless I revise my belief, not come to have a proper understanding of the concept.

In this section, I will argue that it is more likely that Harry lacks the concept All, by arguing that cases of partial understanding, for all practical reasons, do not amount to having a proper grasp of the concept as we do, and can be treated the same way as the case of a complete lack of the concept. First, I will argue that merely employing the concept in observational statements is not sufficient for concept possession, where I will use a line of thought due to Brandom (2000) claiming that inferential capacity is necessary for concept possession. Then, I will argue that a partial understanding is not enough in order to properly individuate a concept. To be more exact, a partial understanding that excludes canonical instances seems, as in the previous section, to charge one with not possessing the required concept, but possibly some other concept. In other words, such a partial understanding underdetermines which concept is possessed. To that end, I will argue that a homophonic translation of such a person's utterances might not be the most charitable option.

4.3.1 *Observational statements and concept possession*

How can we make sense of the idea that Harry has only a partial grasp of the concept? Casalegno (2004) raises an interesting example for conjunction (p. 407), which I will adapt for the universal quantifier. Suppose Mary has a cognitive issue that prevents her from making inferences. She can, nonetheless, competently assert sentences such as *All cookies have chocolate* when I show her a tray of chocolate cookies. She can also assert *Cookie A has chocolate*, and so on, based on simple observations. However, if I first

tell her that *All cookies have chocolate*, then take out a cookie and ask *Does this cookie have chocolate?* she cannot answer if she cannot see it. Thus, one might conclude, following Casalegno, that Mary understands All despite not performing inferences such as (UI). The conclusion, then, is that being able to reason according to a set of rules R is not needed in order to possess a concept C. In the coming paragraphs I will show why we should resist Casalegno's conclusion.

Boghossian (2012, pp. 228-9) tries to resist Casalegno's intuition as well, but I find his argument incomplete. He mentions that Mary's case is "mystifying" (p. 229), but his reasoning is based on not seeing how Mary can assert the right sentence based on observation, but not infer the same conclusion when observation is blocked. His argument is limited to a simple presentation of an opposing intuition, without truly overcoming Casalegno's challenge. I think Boghossian's intuition is on the right track, but fails to explain why we must conclude that Mary lacks the requisite concept. To complete Boghossian's reply, I will use a line of argument due to Brandom (1994, 2000).

Setting aside the apparent unnaturality of both Mary and Harry, we can nonetheless reject Casalegno's conclusion on principled grounds. Only being able to make observational statements, i.e., respond to external stimuli seems to make Mary, as Brandom (2000, pp. 47ff) would put it, no different than a parrot or a thermostat. Then, in order to differentiate responses to stimuli from proper concept possession, inferential ability seems to be necessary. McDowell makes a similar point:

> The ability to produce "correct" colour words in response to inputs to the visual system (an ability possessed, I believe, by some parrots) does not display possession of the relevant concepts if the subject has no comprehension of, for instance, the idea that these responses reflect a sensitivity to a kind of state of affairs in the world (McDowell 1994, p. 12).

On Brandom and McDowell's view, we can very well train a parrot to say *This is red* when presented with red cloths, but this seems no better than conditioning, and we can hardly say that, because the parrot can correctly assert the required sentence, that it possesses the concept Red. Likewise, a thermostat can accurately indicate the temperature, but we can hardly say it possesses the concept Temperature. Correctly responding to observational stimuli, then, is not a sufficient condition for possessing a concept, unless we want to endorse the conclusion that a thermostat possesses the concept Temperature, which seems deeply unnatractive. I believe this conclusion is intuitive on its own, and does not require endorsing any of the views Brandom derives from this lesson. However, it seems that being able to draw at least *some* inferences is then a necessary condition for possessing a concept but,

contra Casalegno, we do not have to endorse a set of formal rules, but only what I have called canonical inferences.

If Mary is incapable of drawing inferences, but correctly responds to stimuli, we can wonder how she managed to learn English in the first place. If she correctly asserts *All cookies have chocolate*, and based on observation also correctly asserts that *Cookie i has chocolate*, but cannot conclude without observation that *Cookie i has chocolate*, then in what way is she different than the parrot who 'speaks' English? Perhaps her expressive power is greater than a parrot simply because she can assert more such observational statements, but this is only a difference of degree, not of kind. Even worse, assuming she is, by Casalegno's definition, unable to perform inferences, then one could wonder if she can sense the incompatibility of her claims. If she says *Cookie i has chocolate* and then that she doesn't know whether *Cookie i has chocolate*, or worse, denies that *Cookie i has chocolate*, then lacking the capacity to perform even the most basic of inferences might make her unable to realise that the statements are incompatible. This is, again, similar to Brandom claiming that a parrot cannot infer that *This is green* and *This is red* are incompatible.

Now, if Harry were unable to perform inferences due to cognitive impairment, then the (AP) would be pointless, as we would be trying to overcome not a logical gap, but a neurological one, and there philosophy is of no help. Padro's assumption is that Harry *is* capable of performing inferences, but is simply unable to adopt a certain rule. In Mary's case, it would be no wonder that adopting (UI) is impossible, since no rule will help her overcome a cognitive disadvantage. In Harry's case, we can set this possibility aside. If Harry has the capacity to perform inferences, then at least in principle, it is *possible* for him to reason in accordance with logical principles. What I aimed to show in the previous paragraphs that it is highly unlikely that someone incapable of performing inferences can be said to possess a concept. However, other scenarios remain open, to which I turn now.

4.3.2 *Partial concept possession and underdetermination*

Supposing Harry is capable of performing inferences, in what way can we say that he might nonetheless possess the required concept? If, by Padro's scenario, Harry has never performed an inference in accordance with (UI), then perhaps Harry accepts some other canonical instances, for example the rule for introducing the universal quantifier: $Px \vDash \forall x Px$, or canonical instances thereof. Then, keeping with the cookies example, he can correctly infer that if I put in front of him all the cookies and tell him there are no other cookies on the tray, he can correctly go from *Cookies a, b, c, ... have chocolate* to *All cookies have chocolate*, but if I ask him later whether, based on his conclusion, *Cookie i has chocolate* follows, he will be clueless.

Depending on how we individuate logical concepts, it could be the case that Harry is using a different quantifier which happens to have the same

introduction rule as the universal quantifier. In that case, a homophonic translation might not be the right move. This would be similar to someone who correctly assents to *Sweden is a democracy*, but also assents to *North Korea is a democracy*. In this case, we would again question the agent's understanding of the concept, even though it might seem he would, at first, have a partial understanding. Then, the question is whether partial understanding allows one to already possess the concept, albeit in a weaker form than us. I think that such a partial understanding underdetermines which concept is possessed, since it is possible that two related concepts can have common instances, for example DEMOCRACY and COUNTRY both have *Sweden* as an element in a canonical instance, e.g., both *Sweden is a democracy* and *Sweden is a country* can be taken as canonical instances.

For example, suppose someone correctly assents, when presented with information that p, q, to *p and q*, and does so with regularity. With the same regularity, when presented with information that *p and q*, they refuse to conclude either p or q. In this case, either the person commits an error, or we are wrongly assuming that they are translating *and* as logical conjunction $\wedge$. Any other information lacking, and assuming a minimal principle of charity, we should question our homophonic translation. In other words, that the way they understand *and* is the same as $\wedge$ (assuming, for the sake of this example, that we understand English *and* as $\wedge$). This case would be similar to Warren's (2015) Tonkers and the trivial Tonklish language. If the connective TONK has the introduction rule of disjunction and elimination rule of conjunction,[29] then, on the face of it, a language including TONK cannot exist. However, this rests on the mistaken assumption that we must translate sentences homophonically, i.e., when someone speaking Tonkish asserts '2+2 = 5', that we have to translate it as the syntactically identical sentence '2+2 = 5', committing them to a mistake. However, as Warren puts it, "if our proposed translation makes those we are interpreting *unaccountably wrong* and *shockingly irrational* then our proposed translation should be rejected" (Warren 2015, p. 8). What I mean to say is that just because someone uses the words 'and', that does not mean they refer to AND.

On the same note, could someone, in this same sense, have a partial understanding of *tonk*? The rules for *tonk* trivialize a language. Thus, supposing someone only assents to the introduction rule for *tonk*, but systematically rejects the elimination rule, would we still be entitled to translate their use of 'tonk' as our TONK, or would it be more proper to translate that person's 'tonk' utterances into utterances containing disjunction, i.e., $\vee$, and thus avoid a homophonic translation? The latter seems to me the right approach, as translating their utterances homophonically will make us conclude that, despite only having a partial understanding, they will be committed to any other sentence they may utter (in virtue of *tonk* trivializing a language), and violate

[29] More exactly, $p \vDash p$ TONK q and p TONK $q \vDash q$, see Prior (1960).

any kind of charity principle.

To push this line of thought further, someone who simply says *yes* to everything would also not count as someone who possesses the required concepts, despite sometimes correctly assenting to canonical instances. Thus, partial understanding seems to fail to individuate any one concept with certainty. More exactly, the claim of the previous paragraphs is that partial understanding, while enough for possessing *some* concept, is not sufficient for possessing the required concept. Thus, if Harry assents to the introduction rule for All, he possesses *a* concept, possibly even All, but this partial understanding is not enough to definitely pick out All from among other possible concepts which have the same introduction rule as All.

Importantly, I am not claiming that such agents would have *none* of the concepts that match the accepted canonical instances. It could very well be the case that one concept is possessed, or even more. The claim is simply that we cannot say with certainty which one it is. For example, someone who assents to the introduction rule for *tonk* but not to the elimination rule could possess either Tonk or Disjunction, but we do not know which one. To give an analogy, it would be similar to a medic making a diagnosis solely based on a single symptom, when there are multiple diseases consistent with the symptom. Nor am I claiming that full mastery is required for concept possession. I only claimed that being able to recognise canonical instances is required to possess the right concept (a similar claim is endorsed by Bengson and Moffett 2007, pp. 42ff). Full mastery might require recognition of non-canonical instances, and disagreement over these instances is allowed. As long as canonical instances are fixed, those who disagree over non-canonical instances are not talking past each other, i.e., their dispute is not merely verbal.

One potential issue with the way I present concept possession is that it might preclude concept learning. I do not think this is the case. For example, children learn target concepts in stages spread over time. Then, must we conclude that they lack the concept during those intermediate stages? I believe so. For example, when a child first learns about dogs, they start to have some understanding of what a dog is, but they will have an incomplete understanding. For example, they could think that any four-legged animal is a dog, and thus upon seeing my cat they will call it a dog as well. It seems intuitively right to say that the child does not yet understand what a dog is, despite having some partial understanding, and being on the right way. They identify the concept Dog with the concept Four-Legged Animal, which is something we differentiate. This is something that will apply in Harry's case as well, meaning that we can treat him just like a child until he learns the concept as needed. Harry will lack the concept in those intermediary stages as well. I am not claiming that there is a clear-cut point when the concept forms, but that it is highly likely that during most of the intermediary stages of learning, Harry will not possess the right concept.

Another worry could be the one raised by Boghossian and Wright (2024), namely that whatever rule you might give me, I will interpret it in terms I already know. For example, in Harry's case, giving him (UI), he can only interpret it in terms of his own practice, and thus the rule cannot extend his practice by incorporating (UI) into it. For example, if a child identifies Dog with Four-Legged Animal, then me telling them *Dogs bark* in order to help them differentiate the concepts, they might then simply think that all Four-Legged Animals bark, instead of limiting the fact to dogs. How can we then help the child make the difference? As mentioned before, examples and forms of non-deductive reasoning play a much greater role than thought when it comes to learning. For example, I can show the kid my cat, and they will only hear her purr or meow, which might make them wonder if she is indeed a dog. Likewise for many other animals. Acceptance of such particular examples is the learning-analogous version of my endorsement of canonical-instances when it comes to understanding, and is indeed a natural view of learning. As my solution in the next chapter will show, Harry has to learn something else before he can understand (UI).

Lastly, one can wonder what is the relationship between partial understanding and misunderstanding *à la* Burge (1979). For example, a patient who claims they have *Arthritis in the thigh* stands corrected by their doctor who tells them that Arthritis can only occur in the joints. These cases of misunderstanding should be distinguished from cases of partial understanding. A case of misunderstanding would be someone thinking the elimination rule for *tonk* is p tonk $q \vDash p$, whereas partial understanding would mean not endorsing an elimination rule at all. While in both cases one would not have a proper grasp of the concept, in cases of misunderstanding one violates a concept's rules, whereas partial understanding means that one fails to endorse a concept's rule. Concept underdetermination, as I conceive it, occurs in cases of partial understanding. For example, if I endorse only $p \vDash p$ tonk q, it is not clear if I have in mind Tonk or Disjunction. Likewise for a patient who could have Arthritis or Arthritis*, where Arthritis* has the same symptoms but occurs in the thigh. The reason they are being corrected by the doctor is, presumably, due to the fact that Arthritis* does not exist as an actual medical condition. If there were such a disease, we could very well imagine the doctor correcting not the location of the pain, but the name of the disease, saying *Do you mean Arthritis*?*. Since Harry does not commit a misunderstanding, we do not have to correct a mistake, and thus we are in a case of underdetermination. It is not the case that Harry uses All in a mistaken manner so we would correct him, saying *Ah, you meant All!* As it is now, we do not know what he has in mind.

Before moving forward, let me summarise the argument of the previous pages:

(1) Possessing a concept requires *some* inferential competence, and thus we cannot say that someone only able to make observational statements possesses the required concept. Otherwise, we would have to endorse

the conclusion that a thermostat possesses the concept Temperature.

(2) Partial understanding underdetermines concept possession. If someone only assents to a subset of canonical instances, then it could be that they possess other concepts, and not the target concept. It could just happen that they use the same word as we do, but have a different concept in mind, such as when a child uses Dog to refer to Four-Legged Animal. In that case, for all intents and purposes, we can treat them no different than someone who lacks the concept.

I am leaving open the path of extending a partial understanding to a full understanding, as I believe this could proceed in similar ways as to what I will describe in the next chapter. In other words, while I think the (AP) can also be solved if Harry has a partial understanding of All, due to the fact that this partial understanding does not lead to concept possession, we can treat cases of partial understanding as cases of lacking the concept entirely. The main difference between the two scenarios is that when we say *all*, then Harry will understand the word in terms of his partial understanding, whereas if he lacks the concept entirely, he will not even be able to pin down some meaning for a word, rendering it practically meaningless.

In virtue of this section, I am ready to modify my initial claim, that of Harry not understanding the rule, as such:

Thesis

Harry does not understand (UI) because he lacks a proper grasp of the concept All used to formulate (UI).

I believe the following analogy can provide some starting intuitions. If we think of concepts as containers, and of rules as specification on how these containers are to be filled, then lacking a container prevents one from filling the container in the first place. What I mean here is that the rule we give Harry is inert in the sense that he does not know what to do with it. Thus, what a practice can afford Harry is the acquisition of such a concept.

To push the idea further, consider the following statement:

**Lasrevinu statements imply each of their *secnatsni.*

I am a trustworthy person, and I assure you this rule will be very useful. Now I tell you that

**Lla objects are black.*

And I ask you if

Is the Colosseum black?

Ignoring similarities to already known words, the above should seem meaningless to you, and of course you cannot conclude what I ask you to conclude.

You are now in Harry's position. You are asked to apply a rule you do not understand, and then blamed for not being able to apply the rule in question.

How could I help you further? Well, one way we learn concepts is by being exposed to examples. You should ask me what is a **Lasrevinu statement*, and if I really want to help you, I should give you examples. I should help you form this foreign concept before I am entitled to blame you for not understanding. This is the moral I will develop in the rest of this chapter. The following last subsection will set the ground for how one can learn concepts.

4.4 Concept learning: brief motivation

In the previous sections, I have argued that Harry does not understand the rule he's being given because he lacks the concept All. I have argued that this diagnostic is not *ad hoc*, i.e., simply to find a way to solve the (AP), but is a principled diagnostic in its own right. However, the (AP) can be raised to the level of concepts, claiming that adopting new concepts is impossible. The next chapter of my thesis will show how we can solve Concept-(AP) in a principled manner. This final section aims to set the stage by discussing concept learning, and distinguishing my solution from Padro's.

A key analogy I will be using is the way children learn number concepts, which then allows them to conform to formal principles which codify the use of numbers. Let's take the case of Hale and Wright's (2001) Neo-Logicist programme in the philosophy of mathematics,[30] the centrepiece of which is Hume's Principle:

> (HP): The number of *F*s = The number of *G*s iff *F*s and *G*s are in a one-to-one correspondence.

This principle is supposed to codify a possible route towards an *a priori* knowledge of numbers. Setting aside such a project's prospects,[31] a natural observation is one made by Wright himself:

> Someone can—and our children surely typically do—first learn the concepts of elementary arithmetic by a grounding in their simple empirical applications [...] [N]o one actually gets their arithmetical knowledge by second-order reasoning from Hume's Principle (Wright 2000, p. 327).

[30] One could take another formal characterisation of the natural numbers, e.g., in terms of Peano Axioms, or another abstraction principle due to Linnebo (2009). While for my proposal the choice of formalisation does not matter, one could lean towards one formalisation over another depending on one's convictions regarding the nature of numbers, see Snyder et al. (2018) for a discussion.

[31] For some criticism, see, for example, Heck (2011) and Nutting (2018).

To link Wright's comment to Harry, it could very well be the case that Harry will not be able to learn the concept All by learning logical rules for its use, as Birman shows, but would rather have to acquire the concept first. In other words, concept acquisition precedes learning formal rules. This seems a promising route, since

> Children do make inferences along the lines of [Hume's Principle] after they have learnt the meaning of the relevant number words (Buijsman 2017, p. 18).

In other words, once the number concepts are acquired, children follow the formal principles which can be said to constitute the concepts, and I believe something similar will hold for Harry. Harry first has to acquire the concept All, and only then will he be able to follow (UI).

Let's take one quote from Kripke himself:

> If someone cannot see the inference from p and if p then q, to q, you will not help him by telling him: 'Look, the following is correct: if p and if p then q are true, then q is true'. This is because if you just add this on as an extra premise, one will still have to use Modus Ponens to get anywhere (Kripke 2024, p. 18).

The idea is the same: telling someone a formal principle is not helpful, but as I have argued so far, we should not therefore conclude that the rule is unadoptable. Instead, if someone does not "see the inference", then the more likely conclusion is that they do not understand the meaning of the concepts involved, and as such they first have to acquire the concepts. In the rest of this chapter, I will discuss how my proposal differs from Birman's.

One might think that my solution essentially reaches the same conclusion as Padro, namely that we require a logical practice before we can adopt logical rules, so why bother with my approach? Let me spell out the differences in more detail. I claim that Harry cannot adopt the logical practice because he does not understand the given rule, in virtue of not understanding the concepts involved in formulating the rule. However, it seems that properly grasping the concept involves accepting its inferential profile, and thus on my way to solve the (AP), I am begging the question against Padro. However, I do not think this is the case. Padro claims that adopting a (basic) logical principle is impossible, assuming one has a grasp of the concepts involved in formulating the principle. Once we reject **Understand**, we need to solve Concept-(AP). I contend that the very way of adopting a concept involves accepting some inferences as canonical, and thus obtaining a partial grasp of its use. Let me phrase my assumption explicitly:

> **Bundle:** Acquiring a new concept enables one to adopt an inferential practice, and acquiring a concept is done through a non-logical practice.

The point is as follows. Learning a rule does not allow one to adopt a logical practice, since one might lack the concepts required in order to grasp the rule, rendering the rule meaningless. Learning a concept, however, is a process that does not proceed in terms of rule-learning, but engages multiple other cognitive capacities, such as analogy, abstraction, etc. Once the learning process has finished (if it ever does, if not then after a certain learning stage), one will have a sufficient grasp of the concept, and will then be able to assess explicit rules for using the concept.[32] This bypasses Padro's (AP) by directly solving the Concept-(AP). Learning concepts *implies* also learning (part) of the concept's inferential profile.

Also a note on terminology. I differentiate between logical and non-logical practices. What I mean by a 'logical' practice is one in which we analyse the concepts we use and try to see inferential connections between concepts. A non-logical practice I take to be an activity in which we engage for some purpose, and which will in turn give rise to certain concepts. Let's take an example to illustrate this better. The concept Truth is analysed as part of the logical practice of philosophers and logicians for various reasons and as part of various debates. However, the concept Truth is acquired by learning a language, not by directly taking part in discussions about the nature of Truth. Then, our practices ground our concepts, which we can then further analyse. We cannot analyse something we have not acquired.

There are three scenarios left possible. One either lacks both the concept and the ability to take part in an inferential practice (e.g., a child lacking the concept of Large Cardinals not being able to take part in discussions about set theory), one lacks the concept but is able to take part in an inferential practice (such as when makes inferential transitions while lacking the concept of Validity), or one both has the concept and can take part in the inferential practice. The possibility that I have been arguing against is one in which an agent has the concept but still cannot take part in an inferential practice, because not assenting to certain canonical instances means one does not understand the concept, and *a fortiori* that one lacks the putative concept. This is what my rejection of **Understand** entails. Having a proper grasp of a concept automatically allows one to take part in an inferential practice.

Where does that leave my proposal? I agree with Padro that before one can adopt a logical principle one must already have a logical practice, but before one can have a logical practice, one must have the concepts involved in that practice. *Contra* Padro, adopting concepts does not proceed in terms of adopting rules or principles for these concepts, thus avoiding circularity. If we would be learning concepts by explicitly learning rules for using concepts, then my solution would be no better than a restatement of the (AP). However,

[32] To take an example, we definitely call a car a vehicle, but we might have doubts regarding a tricycle. However, in order to have doubts as opposed to cluelessness, we must first recognise, as Hart (1961, p. 126) points out, that if anything is a vehicle, a car definitely is one. If I don't know what a vehicle is, not only will I be clueless about the bicycle, but also about *anything*.

I argue that concept learning proceeds differently. Indeed, immersing in a practice will be essential, but it will not be a logical practice at play. Something more basic is needed. Only once you have some kind of concept can we begin to refine it into a precise one. You do not need to accept (UI) before you can start using (UI), but you need to have some idea of the concept all before you can even begin to comprehend the rules for using all.

Finally, just how much can we expect from Harry? Like in most learning environments, some expectations must be met when it comes to the student's cognitive capabilities, otherwise the learning process will be over before it begins. Just how much can be taken away in terms of cognitive power while an agent could still be capable of learning certain concepts seems to me to be more of an empirical question. I have mentioned, however, that some things are nonetheless important, such as the ability to draw some inferences. To still keep the viability of the (AP) alive, we should give Harry as much cognitive strength as a human can be expected to possess, otherwise we could claim that a cognitive deficiency, rather than the unadoptability of a rule, could be the culprit. To that end, I will simply assume that Harry is a human just like us, and that his lack of understanding might come from learning a non-quantificational fragment of English. This is, I think, enough to avoid any problems down the line, and I believe to be a simple, plausible, and minimal explanation of Harry's development.

CHAPTER 5

Concepts and Practice

In the previous chapters we have met Harry who seems unable to adopt the (UI) rule. I have walked through the argument given by Padro to support the conclusion that (UI) is unadoptable, and have shown how it can be resisted. My claim was that Harry did not adopt the rule because he did not understand the rule, which I further refined to the claim that he lacks the required concepts. Therefore, before Harry can begin using (UI), he needs to acquire the concept All, and acquiring a concept does not proceed in terms of learning formal rules, similarly to how our (actual) knowledge of numbers does not come by being explicitly told Hume's Principle. Instead, formal rules codify our competent use of concepts, but acquiring concepts does not proceed in terms of learning these formal rules. This is, I claim, Padro's big oversight. She is trying to use the logical analysis of a concept as the way by which the concept itself is acquired, which, except in highly academic contexts and for specialised/technical concepts, will fail. By blurring this distinction she reaches the false dilemma that either a rule can be adopted by explicitly being told the rule, or it is unadoptable.

In this chapter I will sketch a possible route via which Harry could adopt the concept All. I will point towards a promising route towards reconstructing the way humans might acquire the concept, by focusing on the "functions and genealogy of particular parts of language" (Price 2011, p. 12), and will defend it from potential objections. Due to reasons of space, my proposal will simply lay the groundwork, as a full account would require a thesis of its own. Instead, I will focus on defending the viability of my proposal against potential criticism, as I believe this to be a more achievable goal, and indeed a more pressing one. I will keep for presentation purposes the analogy with

number learning, as I believe the comparison between number learning and Hume's Principle can act as a blueprint to acquiring the concept All and inference rules such as (UI).

Here's an overview of this chapter. I will first discuss the case of number cognition as a case of concept acquisition. Importantly, I highlight that learning numbers does not happen by learning formal rules, and that it is based on some practice, namely counting. Then, I will move on to Harry's case. To teach him the quantifiers as something new, I propose we teach him a game from which the quantifiers get their meaning. I will defend this approach, and show how it manages to solve Harry's predicament.

5.1 Number cognition as a case study

The following observation should be uncontroversial: *We do not learn numbers via formal principles*.[1] What I mean by that is that acquiring the concept Number does not start by being told explicitly the rules of use. If a child has never counted before (setting aside the issue of language), they will not learn to count by being told Hume's Principle. The question, then, is how do we learn numbers if not by following the formal principles? There seems to be a *prima facie* discrepancy between the way numbers are analysed *qua* concepts, and how numbers are acquired in childhood. Number acquisition, it seems, has little to do with formal principles. Children do not start with the formal principles which they then use as starting points for deriving numbers.

Then how do humans form number concepts? In this section I will review one popular line of answering this question. This, I believe, will serve as a blueprint for how Harry could learn a concept, in a similar vein to how humans learn numbers. For the sake of brevity, I will omit many details, and instead focus on the bigger picture and the main philosophical takeaways that I will use later on for Harry.[2] What I mean by 'number concepts' is not systems that allow us to compare cardinalities, but something that allows us to represent *exact* cardinalities. In other words, there is a gap between representing something as being *approximately* Twenty-Four and *exactly* Twenty-four. Whereas the latter is sensitive to small changes, the former is not. Whereas the former system is thought to be innate, i.e., we have an innate capacity for representing (approximate) numerosity (Beck 2015; Dehaene 1997), the latter exact system is thought to be acquired (Carey 2009a).

More exactly, there are two systems which are thought to be innate: the above-mentioned analogue magnitude representation system, which allows us to distinguish numerosity, i.e., approximate quantities, relative size, etc., and

[1] For a different view that claims that children first represent the axioms and then the numbers, see Rips et al. (2008). Even then it is doubtful that they represent the axioms by being *told* the formal principles.

[2] For more thorough overviews, see Buijsman (2017), Carey (2009b), Samuels and Snyder (2024), and vanMarle (2018).

its performance decreases the closer in cardinality two quantities are,[3] and what is usually called the *small number system* (Margolis 2020), which allows us to represent exact quantities with small cardinalities, usually up to four. However, most authors point out that these two systems, while an important precursor, do not suffice for explaining our capacity for representing exact numbers (Carey 2009a; Samuels and Snyder 2024; Sarnecka 2021). Then, it is concluded, numbers go beyond our innate cognitive systems, and are thus acquired. To support this conclusion, animal cognition is also brought as an argument. While many animals share the aforementioned innate systems, only humans are able to represent exact cardinalities above 'four' (see the references in De Cruz 2018, p. 166).

Certain experiments are used in order to support these hypotheses. The most discussed one is called a *GiveN* task. Children are asked to hand a certain number of objects. This task checks whether children actually understand the meaning of the number words, for example that 'one' refers to One. Carey summarises the findings as follows:

> First, children are no numeral-knowers—they cannot even reliably give one object when asked for it. Between 24 and 30 months of age, most English-learning children become "one"-knowers. They can reliably give one object but hand over a random number of objects (always greater than one) when any other numeral in their count list is used in the request. They are in this stage for 6 to 9 months. They then become "two"-knowers (can reliably give one or two objects; chose a random larger number for any other numeral), and then "three"-knowers. Although it is much rarer, "four"-knowers have also been observed. Then, around age 3 ½ on average, English middle-class children become cardinal principle knowers—they work out the numerical meaning of the activity of counting and can now reliably produce sets with the cardinal value of any numeral in their count list (Carey 2009a, pp. 297-8).

Children learn the meaning of the number words in stages, from 'one' to 'four', and then they learn the whole number list in one go. Importantly, there is initially a discrepancy between the number list and the associated meanings. Children are usually able to repeat the count list up to 10, but do not associate the required meanings. The list 'one, two, three, . . . ' is to them no different than, say, 'eeny, meeny, miny, moe'. It is simply a string of sounds which they are able to repeat in the correct order, but which has no other significance to them. Likewise, when learning to count, children can answer correctly a *How many?* question, but they do not realise that the last

[3] In other words, it is easier to distinguish 2 from 8 than it is to distinguish 7 from 8. This system is also thought to be present in non-linguistic animals such as monkeys, rats, birds, or even fish (Carey 2009a, p. 119ff).

number in the count list corresponds to the cardinality of a collection. When they stop counting at 'six', they do not realise that there are six items.

Here are two main takeaways relevant for Harry. Firstly, in the case of number concepts, children need a lot of time to acquire the required concepts, quantified in terms of months or years. Even though the formal principles might be, in some sense, *a priori*, learning to reason in accordance with them is still a time consuming process.[4] Secondly, this process does not proceed in terms of the formal principles themselves. At no point is a child told the Peano Axioms, and it is a sensible assumption that if they were told, they would not be able to understand them. Instead, they acquire the concept via practice, in this case counting. They form the number concepts, reason in accordance with the formal axioms, and even have an informal idea of the axioms themselves (e.g., that adding one to a number is still a number, that there is no largest number, etc.). Basing the formation of a concept in some kind of practice will greatly aid Harry's case. This is the main lesson of this section, which I am now going to apply in the following pages.

Before I proceed, some preliminary qualifications are in order. Firstly, I am not positing an essential connection between one practice and the concept it generates. While the practice of counting gives rise to the concept Number, I am not claiming that it is the *only* such practice. I do not know how such an alternative practice would look like, but I am not going to project my lack of imagination. Secondly, and related to the previous point, I do not claim that this schema, of acquisition via practice, generalises to all kinds of concepts, but only to *basic* ones. For example, I do not think there is a special practice, separate from other practices, which gives rise to the concept Blue Triangle. Nevertheless, for basic, primitive concepts, I believe there is some stronger connection between the practice and the concept, at the very least when it comes to the acquisition of the concept.[5]

5.2 A sketch for Harry: 'searching-and-finding' games

In the previous section I presented one plausible account of how children learn the number concepts, and the main resulting relationships between the different notions are summarised in the left-diagram in Figure 5.1. My aim in the rest of this chapter is to defend an analogous picture for the concept All, as in the right-hand side of the picture. Then, I believe the ground is set in order to sketch a promising route via which Harry could acquire the

[4] This might point towards a view such as the one proposed by C. Jenkins (2008), where "experience grounds our concepts (which is not the same as supplying evidence for any proposition), and then mere conceptual examination enables us to learn arithmetical truth" (p. 4).

[5] Berto and Restall (2019) point out a tentative list of such concepts. Focusing on negation, they mention that it can be grounded in the *primitive* (in)compatibility relation (see also Berto 2015, §3).

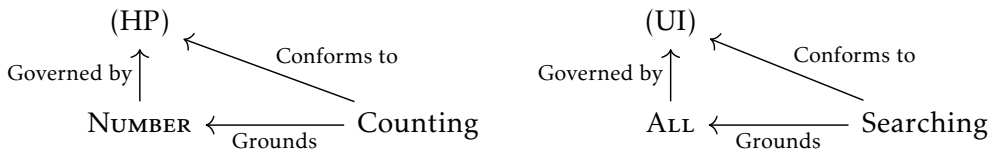

Figure 5.1 – The analogy between Number and All.

required concept. My solution takes a page out of the pragmatists' book, asking *Which practice gives rise to quantification?* While, as mentioned, the focus will be on defending the approach rather than providing a full-blown reconstruction, I believe an initial sketch will provide enough in order to solve Harry's predicaments. My solution will answer, in a plausible way, how Harry could come to adopt a concept, and leaves room for discussing how a concept comes to be part of a community's resources. In this sense, my approach can naturally be extended to cases beyond the individual one of immersing someone in an already existing practice, such as Harry. I will only discuss the individual case, and simply indicate how I think this could be extended into a proper account.

5.2.1 *Human activities, logical tools*

As hinted in the previous section, I follow the assumption that activities are deeply intertwined with our conceptual practices. The natural question, then, is what practice might give rise to quantificational tools. This is where I mentioned that Finley's (2024) solution falls short of satisfactory, and on which I aim to improve in this chapter. I believe an answer due to Hintikka (1973) is on the right path. In Harry's case, we are met with the challenge of teaching someone quantification as something radically new. We assumed that Harry lacks quantificational resources, and the usual method of explicitly stating the rule will fail for Harry. The approach I want to propose is based an observation made by Hintikka, who aimed to give a formulation of the meaning of quantifiers based on Wittgenstein's idea of language-games:

> [T]he meaning of quantifiers lies in their role in guiding the processes ('games') of verifying (in principle) the sentences of our language. The primary meaning of quantifiers, in brief, is their use in the language-games of seeking and finding (Hintikka 1973, p. 103).
>
> From our point of view, to convey the meaning of quantifiers to someone as a radically new idea is to teach him to play the language-games that go together with quantifiers, and to teach him the relation of the different quantifying expressions to these games (Hintikka 1973, p. 58, fn. 12).

Hintikka's proposal laid the groundwork for what came to be known as Game-Theoretical Semantics (for an overview see Hintikka and Sandu 1997). The idea is, roughly, that evaluation of a statement is treated as a two-player game, between Myself and Nature. A sentence is declared true iff Myself has a winning strategy, and false otherwise. The rules are fairly straightforward. Whenever we encounter an existential quantifier, Myself gets to choose the object, and when we encounter an universal quantifier Nature gets to choose it. The game ends when we reach an atomic formula.[6] In the following pages I will exemplify this approach, and defend it as a possible way via which Harry could adopt the required concept.[7]

Let's take an example to see how this works. Suppose we want to evaluate the following formula:

$$\forall x \exists y \varphi(x, y)$$

Nature gets to choose an object in the first round, say a and then Myself has to reply with an object b such that $\varphi(a, b)$ is true. If Myself can always find such a b, then the sentence is true. However, if Nature selects an a for which Myself cannot find a b, then the sentence is false.

Why does Hintikka appeal to 'seeking-and-finding' games? Hintikka puts forward his approach as a "transcendental exposition of quantifiers" (p. 121), which he grounds in his reading of Kant. I will not discuss this aspect of his proposal, as I believe it can be used independently as a learning method for Harry. Nonetheless, without relying on any strong evidence, it seems intuitive that the 'seeking and finding' game[8] is a centrepiece of current human activity. We are always searching for something: food in the forest, a warm home, counterexamples to arguments, a better future. The idea of connecting this behaviour to the rules of logic seems to me, with Harry in mind, appealing. Intuitively, if we are entertaining a quantified statement, it only makes sense to try and think of what objects could validate or invalidate the statement. For an existential, it's enough to find one object, for a universal, we should consider counterexamples, and again, it's enough to find one.

Why should we use Hintikka's approach to the meaning of the quantifiers? To fully appreciate his approach, it will be useful to emphasize why the usual logic-teaching methods will fail for Harry. Let us take a first pass at how the meaning of quantifiers is usually taught in a logic class. The way I have been

[6] For simplicity, I will not touch upon infinite games. Likewise, for the purpose of my argument it will be enough to discuss the quantifiers. For a more thorough presentation, see Hintikka (1973, Ch. 5).

[7] It is highly likely that the acquired concept would be a practical one, since it acts as an incentive to produce a certain kind of behaviour. While I believe it would be natural to cash out the argument in terms of know-how and productive reasoning (Pavese 2021a), for reasons of space I will not pursue this line of thought.

[8] Where 'seeking' and 'finding' have a very broad meaning, including, e.g., constructing, imagining a situation, etc. The right interpretation will depend on the topic of the statement.

taught, the way I have been teaching, and I imagine the way most students in their first logic class learn the meaning of the first-order quantifiers is via the usual Tarskian inductive definition:

> '$\forall xPx$' is true iff <u>for all</u> objects t of the domain, 'Pt' is true

The problem with this definition for Harry is that it presupposes competence with quantifiers in a meta-language (here, English). In other words, in order to learn the meaning of '$\forall$',[9] one has to know what *for all* means. The point is, in a somewhat humorous way, put forward succinctly by Girard: "The rules of logic have been given to us by Tarski, which in turn got them from Mr. Metatarski" (Girard 1999, p. 219). The straightforward Tarskian way of conveying the meaning of the quantifiers will fail for Harry since, as I have argued so far, Harry does not understand what ALL[10] means, and *a fortiori* he doesn't understand quantification in English. The challenge, then, is to teach someone quantification as something new.

This challenge does not arise just for truth-conditional semantics, but for inferentialist semantics as well (see, for example, J. Fodor and Lepore 2007, p. 682). Thus, it seems likely that the traditional ways of conveying the meaning of the quantifier, either via rules or via Tarskian definitions, will be of no help to Harry. The rule-based way, as shown in the previous chapter, is doomed to fail since Harry lacks the concept used to state the rule. Indeed, quantification is presupposed when stating the rule:

> *Universal statements imply <u>each</u> of their instances*

On the other hand, Tarskian definitions define the object-language quantifiers using meta-language quantifiers, and thus either leads to a regress, since the lack of quantification leads to Harry not understanding the meta-language quantifiers either, requiring meta-meta-language quantifiers and so on. Importantly, I am assuming that Harry lacks quantification capacities entirely, in order to prevent the trick of defining the universal quantifier in terms of the existential one.[11] The real challenge, and the more interesting one, is to teach Harry quantification as something radically new. This is where Hintikka's 'seeking-and-finding' games come in. They provide the background against which the use of quantifiers becomes intelligible for Harry.

[9] We do not even need the logical formalism to convey the point. We could just as well say that for 'Everything is *P*' to be true, it must be the case that every object is *P*. Here English is both the object language and the meta-language, but we are still presupposing competence with quantification.

[10] As noted by Wierzbicka (1996, p. 47), since the concept ALL is hinted at in many English phrases such as *at all*, *already*, *always*, etc., it could very well be the case that the lack of ALL might reverberate, and lead to one also not understanding the mentioned structures.

[11] Birman (2024, pp. 50-1) makes a similar remark.

It should be straightforward to see that the described game corresponds to the usual meaning of the quantifiers.[12] Whether this game-theoretical approach is the right approach to the meaning of logical terms could and has been debated (for a critique, see Hodges 2001; for a defence, see Pietarinen 2006). However, it might prove especially useful in order to teach Harry quantification, and this claim might benefit from the anecdotal fact that Hintikka's method forms the basis for one in-use teaching method (Barker-Plummer et al. 2007). As we have seen, the usual direct method of conveying the meaning of quantifiers fails for Harry. However, this indirect method of 'seeking-and-finding' games might provide a helpful alternative. Indeed, I believe that only such an indirect method, together with some analogical thinking on Harry's part, could solve his predicament.

My appeal to Hintikka's games should be seen as trying to find an analogue to something akin to the *GiveN* task in the case of number cognition. Since we cannot split someone's mind open and look for the concept in there, we are looking for the next best thing, namely the manifestation of the concept's possession in action. One way of specifying this, a way which I endorse, is the following:

> Thus, for a person to grasp the universal quantifier is for that person to be sensitive herself to the possibility of counter-examples to some claim. If a universally quantified proposition is adopted tentatively, if someone, as it were, asks herself whether all Fs are G, then the paradigm manifestation of this is another search, a search for something that meets the description 'an F that is not G' (Evnine 2008, pp. 50-1).

Thus, to question the truth of a universal statement means that one is capable of searching for a counterexample, and recognise the found object as a counterexample. Importantly, I am not claiming that enabling this disposition is the main role of the quantifier, just as manifesting the ability to give N objects is not the main role of the number concepts. Thus, it is important to keep separate the counterexample-sensitivity as a manifestation condition of the concept All, and the point of the concept, on which, for the purpose of this thesis, I will remain agnostic.

Now, why should we think that this connection, between seeking-behaviour and concept possession, must hold? Evnine, cited above, responds to Cherniak's argument that there must be some inferential abilities constitutive of thought. On Cherniak's view:

> [S]o far as rationality is concerned, an agent must only be an adequate logician, not "the right kind" of logician; beyond empirical

[12] And indeed, this is a formal result: A sentence is true iff Myself has a winning strategy, see Hintikka (1973, Ch. 5) and Tennant (1987, pp. 173ff)

> generalizations about human psychology, there is no particular right kind of logician (Cherniak 1986, p. 39).

In other words, there is no fixed set of inferences that an agent must be able to perform, but this will depend on a creature's psychological construction:

> Generalizations about which are the obvious equivalences are contingent cognitive psychology; a deviant creature, or merely a human being with training in formal logic, can violate these generalizations and still be a logically competent agent [...] An agent must only be an adequate logician; beyond generalizations about human psychology, there seems to be no transcendentally "right" kind of logician. (pp. 47-8).

For example, while we might find the inference from the Peano Axioms to Fermat's Last Theorem deeply unobvious, another community with a different psychological build-up might find it extremely intuitive (re-using an example, this is the case of the Mathoids in Berry 2013). This is to say then that whatever we find constitutive of a concept, other communities might not find so constitutive. Then, why should we think that the relationship between the concept All and Hintikka's game should hold?

However, I do not think Cherniak's argument is a worry for this connection. What we are teaching Harry to do is not to perform some inferences, such as (UI). Indeed, if we were to be teaching Harry (UI) and the introduction rule, then Cherniak's worries might hold some water. However, we are teaching him an activity, and not some particular set of constitutive inferences. Rules such as (UI) will be a natural consequence of this set-up, but so will other rules, such as the ω-rule. We might find the (UI) rule primitively obvious, but it could very well be that some alternative agent who performs well in infinite cases but poorly in finite cases will find the ω-rule much more intuitive. Nonetheless, that alternative agent will still follow the exact same rules of the quantifiers, and still engage in 'seeking-and-finding' games. The only difference will be related to the domain of search, or other such things, not to the action itself. The behavioural pattern is identical. More exactly, what we are teaching Harry is an activity in which the concept of the quantifier is embedded in a natural way. In learning to play the activity, one will come to also learn to use the quantifiers, in a similar fashion to how in learning one's first language will automatically make one follow the rules of grammar without explicitly being told the rules. Cherniak's point can then be naturally seen as the claim that which rules are taken to be constitutive of a concept might depend on the community's psychological profile, a claim which I am happy to concede.

From this discussion, since an activity takes center stage and not a certain constitutive rule, the Hintikka-games can be easily moulded into giving rise

to different quantifiers, which I take to be a virtue of the given approach. What I mean by that is that by learning to play some variant of a seeking-and-finding game one will be able to adopt, say, the classical quantifier, the intuitionistic quantifier, and so on. The activity prompted by a quantifier will still be one of seeking-and-finding, but we modify what might count as a valid counterexample. For example, when trying to verify a claim such as *Everything exists*, whether we count Pegasus as being a valid counterexample is irrelevant to the core of the quantifier. We still search for something, but we change what we accept as a legal move, so to speak. I take this to be a virtue since we do not make the way of learning the meaning of the quantifiers *a priori* favour a certain logical interpretation. Therefore, my proposal is agnostic when it comes to logical pluralism vs. monism, and as such can be endorsed without being said to favour any one logical system (and the existence of Hintikka-style semantics for non-classical logics should provide some strength to this claim, see for example Başkent 2016; Hartonas 2019; Majer and Fermüller 2018). For simplicity, I will only consider the classical quantifier.

Now, I want to discuss a bit more the point of the concept All and the point of the seeking-games. For some concepts, at least for basic logical ones, it seems plausible that one of their points is to guide inferences, and that to master these concepts one needs to (at least partially) grasp the point of the concept.[13] To use one of Dummett's analogies, it is one thing to learn the rules of chess, but another thing to learn that the *point* of the game is to checkmate the opponent's King (Dummett 1973, pp. 296f). This is also in-line with a criticism made by Tennant regarding Hintikka's approach:

> Merely being on the lookout for a certain kind of object, without assurance that every object in the domain will eventually be encountered; or casting around blindly for, or randomly grabbing, are not activities in which one diagnoses the execution of a strategy calling for careful and methodical inspection and selection (Tennant 1987, p. 176).

In both cases, that of chess and of Hintikka's games, the activity has a goal towards which the moves are directed. As expected, I believe that this analogy can be applied to Harry's case. Teaching him the (UI) rule and nothing else is similar to teaching him the rules of chess, without teaching him what the point of the game is. In that sense, he has failed to grasp the purpose of the rule, and implicitly will fail to recognise when to apply the (UI) rule, just as a chess player will move the pieces aimlessly if they do not know that they need to checkmate the King, or even worse, move them in a way that will make them lose if they are, say, under the impression that the purpose is forcing the

[13] For more on the point of concepts, see Queloz (2019). For example, the point of negation is sometimes said to be that of signalling incompatibilities, see Price (1990).

opponent to checkmate them. Teaching someone the point of a concept places the rules in some context:

> the rules present to us as important and discernible realities only within the perspective of practices that have an independent appeal (Pettit 2023, p. 143).

How far the appeal of the seeking-and-finding games goes is beyond the scope of this thesis. I hope their intuitive connection with how most of us do logic is clear enough, but if one is not convinced, then I hope one can at the very least grant the weaker point that it is *one* plausible way to learn the meaning of the quantifiers, and since there is a proven equivalence between the games and the usual model-theoretic semantics, this point should be granted. It could very well be that other activities might lead to the learning of quantifiers, but as Hintikka notes

> [...] [I]t seems to me that [the game of seeking-and-finding] is by far the most important kind of language-game in which [the quantifiers] can occur (Hintikka 1973, p. 59).

Going back to the question at hand, what could be the point of the semantic games? I am not asking why we engage in these games, that is a different matter. I am asking more in the sense of Dummett's chess analogy mentioned above. For lack of a better word, it is *winning* the game. This might seem uninformative, so let me discuss it a bit further. Checkmating the opponent's King is also winning, but clearly the winning conditions are different in the case of Hintikka's games. To be more precise, winning here means verifying the sentence.[14] We want what we assert to be true, and what we doubt to be challengeable. Understanding the quantifiers allows one to competently defend a challenge to one's assertion, and in turn to challenge someone else's assertion. And just like in most cases, winning once does not guarantee winning always, just like finding a true instance does not guarantee general truth. One could very well grasp the goal of a game but be a lousy player. For example, I understand that the goal of chess is to checkmate your opponent, but I'm still bad at playing it. In that sense, having a goal directs your actions, but having a strategy is not required for engaging in play.

Another worry could be that the turn-based method of evaluating the statement might prove troublesome. Being able to evaluate statements is something that usually happens at the individual level, as we will not always find an opponent against whom we can play the game. This is, for example, Hodges' (2001) main criticism of game-semantics. I think, however, that in Harry's case this oversimplifies Harry's overall capabilities. After all, simulating behaviour is a capability most complex animals have. However,

[14] Not to be confused with truth being equivalent to having a *winning strategy*.

there is a difference between mere imitation and actually simulating someone else's role in an activity. To overcome this worry, I will take onboard a recent proposal due to Paulson (2024) as to what internalizing a rule amounts to. The main idea is that a student can *toggle* between the role of the teacher and the role of the student when performing an activity.

When learning an activity, a child might at first simply imitate what they are seeing, and react to certain corrective behaviour from an adult. However, as the child's cognitive capacities develop, a new form of behaviour is enacted, namely what is usually called *self-regulation* (Tomasello 1999). Tomasello sums it up as follows:

> as the adult regulates the child's behavior in some cognitive task or behavior, the child attempts to comprehend that regulation from the adult's point of view (to simulate the adult's perspective). And then, in many cases, the child later reenacts the adult's instructions overtly in regulating her own behavior in that same or a similar situation (Tomasello 1999, p. 192).

Therefore, during cognitive development we not only become capable of imitating others' behaviour, but also try to actively comprehend their behaviour from their point of view, which, later on, we can simulate without their explicit instruction, or even presence. Internalizing a rule, then, is the capacity to self-instruct in this fashion. Since I mentioned that I take Harry to have the same cognitive capabilities as the average human, I see no reason to think he would not be capable of this behaviour.

Let me summarise the main takeaway of this section. What I want to disentangle, and highlight, is that it is one thing for a practice to give rise to a concept, and another for a concept to be governed by certain rules. Remembering the diagram at the beginning of the chapter, I keep the practices, the concepts, and the rules separate. What I am claiming here is that it is one thing to learn the practice and thus acquire the concept, and another thing entirely to recognise that, for an *already* known practice, some rule holds for a concept used in that practice.[15] One does not first learn to count, and then conform to (HP), but in learning to count, one will *automatically* conform to (HP). Then, learning how to search for something is one activity which will give rise to the concept All, which in turn is governed by rules such as (UI). However, as mentioned, one will not understand the rule without obtaining the concept from some practice. Then, Harry has to follow two steps. First, he has to learn to play the searching-game, and secondly, he will have to recognise that (UI) holds for the concept All acquired from the searching-games. The adoption problem for rules is entirely bypassed, as the acquisition of the concept will automatically make one conform to some rules which are relevant to the concept. The way I think is as follows:

[15] Remember the analogy with (HP), counting and the number concepts.

- The (AP) does not arise for rules, since Harry will not understand the rule (as argued in the previous chapter).
- The Concept-(AP) can be solved, since it is possible to acquire primitive concepts, but this is not done through adopting rules.
- The concept All will be acquired by teaching Harry a practice which uses it, similarly as to how the concept Number arises from the practice of counting.

Would I go as far as to say that "the true basis of the logic of existence and universality lies in the human activities of seeking and finding" (Hintikka 1983, p. 33)? Perhaps not, but I do not have to go that far. It is one thing to claim that the quantifiers get their meaning from these activities, and another to claim that they are a good learning method. For the time being, I will remain agnostic on the former, and simply claim the latter. Supposing one is convinced of the connection between the activity of 'seeking-and-finding' and the meaning of the quantifiers, what next? Let's assume Harry has learned how to play the game. The next step is to check whether he can adopt the (UI) rule, now that he has formed the required concept. This is the focus of the next subsection

5.2.2 *Implicit learning and the dissolution of the adoption problem*

Supposing Harry has learned how to play these games, will he be able to follow and adopt (UI)? Suppose I give him a statement, say *Everything is self-identical*. What does stipulating that it is true do for Harry? And, importantly, how will he be able to recognise the (UI) rule as a good rule within the context of the game? This is the focus of this subsection.

First, the game rules are given for connectives. Since (UI) is a rule of inference, and not a connective, we need the notion of consequence in terms of game-semantics. Hand (1988) provides some guidance in this respect. Starting with the observation that the orthodox criterion for inference rules is preservation of truth, translating into game theoretical terms means that the existence of a winning strategy for all the premisses translates into the existence of a winning strategy for the conclusion. I suggest that for Harry to conform to the (UI) rule, the following should hold:

> Where a universally quantified proposition is accepted, this is manifested in the tendency to select, without checking, an F if one is after something that is G. Its acceptance is freedom from the expectation that a given F may fail to be G (Evnine 2008, p. 51).

Then, this is again manifested in a disposition. However, this disposition is different than the one presented by Devitt and Roberts (2024). In their case, the disposition arises by simple conditioning to infer according to the (UI) rule, whereas here the (UI) rule arises as a natural consequence of a practice,

that of searching. In other words, simple conditioning will not allow for the formation of the concept, but taking part in a practice will. Why? Because within a practice, the concept will make sense, meaning that it has a clear guiding role, whereas in the first case this will not be the case. The end result will be the same in both cases, namely the disposition to infer according to (UI), yet the way this disposition is acquired is different. While in the first case it is conditioning, in the second it will be the result of possessing a concept.[16]

To illustrate this further, consider the case of counting. Without the practice of counting, addition is defined in a purely mechanistic fashion as symbol manipulation. However, once we have the practice of counting, addition will make sense, since a child can think as follows: "If I put together 2 things with 3 more things, then if I count them I get 5 things, and it so happens that '2+3=5'". Then, addition will make sense by again correlating it with the practice of 'putting-things-together'. A similar case can be made for subtraction (staying within the case of positive integers), and other arithmetical operations.

Learning to use the (UI) rule might proceed, initially, by simple association. In simple action terms, he might observe that whenever we stipulate a universal as true, and he picks something at random, that random object will satisfy the conditions. Thus, he will come to accept the (UI) rule as something that holds, in the sense of Williamson (2024) who says that "accepting principles of mathematics or logic outside a formal language often involves messily defeasible dispositions not utterly unlike those involved in accepting theories elsewhere. Such dispositions are not normally immutable. We can consider some cases of changing one's dispositions to reason in various ways in mathematics and logic" (p. 424). Again, this is given in terms of recognizing certain patterns in the activity, but importantly enough this recognition comes within a certain practice, that of searching for things. Harry, once he understands the meaning of the quantifiers as prompts for searching, he will also be in a position to understand the (UI) rule. Nonetheless, would this count as *adopting* (UI) in Birman's sense?

Let's recall the three conditions laid out by Birman on adoption. I will comment how my approach interacts with Birman's original formulation, and highlight why I believe the (AP) can be overcome.

(1) One does not yet reason in accordance with the rule. Here, we must be careful as to what *reasoning in accordance with a rule* amounts to. If we take a strict interpretation of not directly performing an inference according with (UI), then there is a gap between learning the rules of the game, on the one hand, and learning (UI), on the other. (UI) is

[16] Someone might worry that we are nonetheless conditioning Harry into playing the seeking-game, and as such at the end there is conditioning nonetheless. However, I think this is misguided. The analogy I take is one with learning one's first language and then following rules expressed in that language. In order for someone to understand the rule, they must first learn the language, and learning one's first language is not a matter of learning rules.

nowhere to be found in the rules of the game, but due to the setup of the game, it will naturally hold. On this view, Harry does not yet reason *according* to (UI), but will conform to it.
However, I find it more natural to say that Harry *will* reason in accordance with the (UI) rule once he learns to play the game. This is because the rule is implicit in the game, in a similar vein as to how (HP) is implicit in the practice of counting. I find it an odd conclusion to say that in order to conform to (HP), a kid needs to not only learn how to count, but also to be told the principle implicit in the act of counting. Conforming to implicit rules is still a case of reasoning in accordance with a rule.

(2) One comes to view the rule as a good rule. In the original scenario, we are simply told that Harry is supposed to trust us that the (UI) rule is a good rule. However, this acceptance seems simply to be an external one, and not one due to understanding the rule. In the case of the 'seeking-and-finding' games, Harry will have come to learn an activity in which the quantifiers are embedded in a natural way, and thus will also be in a position to recognise whether the (UI) rule is a good rule or not.

(3) It is in virtue of accepting the rule that one comes to reason in accordance with the rule. Once Harry learns the Hintikka-games, he will then be in a position to recognise the (UI) rule as a rule that holds *as a consequence* of how the games are set-up, and thus also be in a position to start using it whenever he thinks it might be useful. The (UI) rule can be seen as a shortcut when it comes to verifying certain sentences. For example, if one has already verified a universal statement as true, then one will not have to play again, one can simply instantiate the statement. Simply put, Harry will be in a position to both understand that the (UI) rule holds, and thus accept it as a good rule, *and* to apply it by thinking in terms of searching activities and winning strategies.

Nonetheless, the apparent problem here is that we are no longer in a position to adopt a rule, since condition (1) will not be met. Harry *will* reason in accordance with the rule once the required concept is learned. This should not be surprising. The whole point of the previous chapter was that understanding a concept implies finding some canonical instances compelling. As such, once the concept is acquired, Harry will *also* come to follow the rules that govern that concept. This is why I believe, as mentioned above, that the (AP) is a false problem when understood in terms of concepts. You do not first acquire the concept, and then adopt the rule, but in acquiring the concept you automatically come to follow some rule. You cannot be said to understand what All means unless you are prepared to admit that from $\forall xPx$, Pa also follows.

Of course, one could reply that Harry might arrive at the (UI) rule outside

such a practice. After all, he can also associate true sentences, as in observing that whenever a universal sentence is true, its instances are also true. However, I think this is too hasty. Indeed, we can always associate things for whatever reason, but I think that in the case of the behaviour of seeking, this connection is much stronger. Whereas in the first case he would merely observe the association,[17] in the case of the seeking-games he would be in a position to understand it in a proper, quasi-causal way. What I mean by that is that in the associative case there is usually no explanatory link between two events, but in the case of the games, the association can be explained, as the rule will *make sense* in the context of the games. Acquiring dispositions through mere conditioning will not give rise to a concept. As Weiskopf mentions, "[j]ust observing a set of correlated properties won't produce a new concept if a person doesn't think that those correlations have a common causal source" (Weiskopf 2008, p. 374). The background of these games provide the common source required for coining the new concept.

Let me say a bit more on the relationship between concept acquisition and constitutive rules. As Harry's case shows, one cannot acquire a concept by being told rules for its use, not even by being told its constitutive rules for use. Telling Harry that *Universal statements imply each of their instances* is of no help if he doesn't know what this rule says, even though it is a rule constitutive of universal statements. Something else is needed to tell him what a universal statement is, what an instance is, and so on. A rule for using something, importantly, cannot get us beyond what we already grasp (Boghossian and Wright 2024), a different kind of method is required. If I tell someone who has never seen the colour green before to throw garbage into the green bin, should I be surprised if they don't know where to throw it?

Then, what is needed is a way for Harry to know what a universal statement is. Importantly, this is different from understanding the concept All. There are two steps needed. First, Harry will need to recognise universal statements, which can proceed in purely syntactic fashion (Besson 2019). Then, from here, he will need to learn how to use these statements. Recalling from the previous chapter, there is a difference between simply hearing a word and then forming a thought associated with that word (e.g., when a kid hears that '2 + 2 = 4' without having encountered addition beforehand), and having a proper grasp of the concept (e.g., a kid actually being competent with addition). As a reminder, I was skeptical of Besson's approach as it stops short of helping Harry get the point of the concept All, and as such he will fail to acquire the concept, and in turn will fail to understand universal statements beyond a purely syntactic device. In other words, his dispositions will again be nothing more than the result of simple correlations, without grasping the concept used. The universal quantifier must be grasped as a prompt for a

[17] Besson and Hattiangadi (m.s.) make a similar remark, distinguishing between inference and mere transition in thought.

certain activity, not simply as a symbol.

Finally, how will Harry recognise that the (UI) rule holds, and understand it in terms of the game of seeking-and-finding? Let's assume that he is able to recognise universal statements, e.g., by pointing them out to him, in a syntactic manner as suggested by Besson (2019).[18] He might first come to observe that whenever he wins the universal statement, he also wins the instance, and thus form an association between these two events.[19] Until now it seems that this association is no better than the cases of conditioning I criticized above, but the next step comes in. He will be able to recognise that this association holds *because* the same strategy he applies in the universal statement also works for the particular instance. In other words, once he has a winning strategy for the universal statement, he will also have a winning strategy for the instance. Why? Since Nature makes the first move, he will have an identical strategy in both cases, as his first move will not change. Therefore, he will be able to appreciate the connection between the two situations in a way that simple conditioning would not enable him. Recalling Weiskopf's quote above, he will form the concept All because he thinks that the two correlations have a common cause, namely the sameness of the winning strategy.

5.3 Objections

In the rest of this chapter, I will pre-emptively discuss some potential objections to the approach sketched in the previous pages. More exactly, I will focus on two aspects. First of all, I will briefly discuss the rationality of concept acquisition, and then the problem of rationalization and rule-following, comparing how my solution fares against Devitt and Roberts (2024). Secondly, I will investigate whether one needs competence with the (UI) rule in order to learn to play the seeking-game.

5.3.1 *Rationality and rationalization*

A quick reply by someone sympathetic to Birman could run along the following lines:

> **Objection:** We want adoption to be a rational process, namely that one begins to use the rule *in virtue of* accepting it. Learning the rule indirectly via some other practice seems to not be a case of adoption, because one come to follow the rules *not because* one accepts them.

First, I want to disentangle some similar uses of rationality that are at play in the cases of concept acquisition and rule adoption. I think we can

[18] See also Sabo (2014).

[19] He could form perhaps an associative concept as suggested by Margolis and Laurence (2011).

distinguish the following two relevant senses in which 'rationality' is used. There is first the problem of the rationality of adopting a rule or a concept, and then there is the issue of rationalization, namely following a rule in virtue of accepting it. I will quickly discuss the first one, and then move on to the latter.

Since my solution is somewhat similar to the one put forward by Devitt and Roberts (2024), some of their replies can be used here as well. However, a more pressing worry arises from my presentation, namely can concept acquisition be rational? Is learning a new concept a rational process? The question of the rationality of concept acquisition seems hard essentially because, from the point of view of the person lacking the concept, it might be hard, if not straight out impossible, to know beforehand if the concept they are going to acquire will be a good one.

Indeed, this seems to be simply a special case of the general issue of incommensurability (Kuhn 1982). However, I want to set aside the issue of the rationality of concept acquisition for the purpose of this thesis. The question is parallel to the one that could be posed in the case of the usual adoption scenario. Is Harry rational for accepting the (UI) rule? A classical logician will say 'yes', whereas a free logician might disagree. However, Harry's own input is not considered, and Birman never considers the issue of why Harry should accept the principle, except that we tell him it is a good principle. The focus, in the classic scenario as well as in the case of where we want Harry to acquire the concept, is on the consequences of accepting a principle/acquiring a concept. While the issue of whether Harry himself is rational for accepting the principle, or agreeing to undergo the process of concept acquisition, is important, I will set it aside for the purpose of this thesis.[20]

From Birman's perspective, the main issue seems to be that Harry will not follow the (UI) rule because he accepts it, but for other reasons, and as such this would not count as adoption in her terminology:

> [T]he possibility of developing the practice by training or psychological conditioning is excluded. There should be no doubt that Harry's acceptance of the principle is the reason and explanation for his inferring in accordance with it (Birman 2024, p. 40).

However, I think there is a difference between what Birman has in mind when she mentions training, and what I propose. I believe her account of training is more aimed at the kind discussed by Devitt and Roberts (2024), where Harry is conditioned to apply the (UI) rule. Devitt and Roberts do not mention how the training process would look like, but presumably it would

[20]The problem can easily be cashed out in terms of framework-choice or paradigm-choice issues, and then discuss the rationality of paradigm-shifts (Kuhn 1982) or of choosing a framework (Steinberger 2015). Discussing this problem, let alone making some dents in it, would require a thesis of its own, or even more.

consist of many cases where the coach would reward Harry for performing well, and penalize for performing poorly. Then, Harry will conform to the (UI) rule without having actually grasped it properly. He is only following it because that's what he was conditioned to do.

My proposal, however, goes deeper by having the (UI) rule as a consequence of the way we use the quantifiers. Quantifiers, on my proposal, are a prompt for a certain kind of activity, which is different from simply inferring. While there is a sort of training involved, it is not the same one as proposed by Devitt and Roberts. We are not teaching Harry to simply use the (UI), but actually teaching him the meaning of the quantifiers. We are teaching him the meaning of All, and the (UI) rule will be a natural consequence of the way he learns to use All.

To better illustrate the difference, consider the following example. Suppose my friend does not speak French, and asks me for advice. I tell him the following:

> Whenever you pass by a French person, say *Comment ça va?*

He does not know what that means, and I can act as a coach, presumably telling him when to say it or not, who's French and so on. Telling him this and training him to do so is analogous to Devitt and Roberts' training method. Now consider the case where my friend actually learns French, and comes to understand that *Comment ça va?* means *How's it going?* In that case, he will come to follow the rule as a consequence of learning French.[21] While the training was focused on something else, he came to follow the rule because he understood the role the rule plays in the overall practice. Further, the practice can have some intuitive appeal on its own, whereas the conditioning seemed to be a conditioning for its own sake.

Birman's worry is, I contend, misplaced. Once Harry will learn to play the games, he will follow the (UI) rule *not* because we are conditioning him to follow the rule, but because the rule is a simple consequence of the way the game is being set up. He is thus in a very simple position to accept the rule, and to immediately start following it once he understands the meaning of the quantifiers. Indeed, the training comes in the form of learning the rules of the game, but that should not worry us.

Now, while Harry does not follow the rule because he accepts it, we can nonetheless say that the rule somehow guides his actions, since the rule is constitutive of the practice of engaging in 'seeking-and-finding' games. This might be a roundabout way of somehow being guided by the rule, by being immersed in a practice which has the rule as constitutive. For example, while we may not explicitly accept the rules of grammar, they nonetheless are constitutive of the practice of speaking. In other words, when we speak, we

[21] Assuming manners and conversational conventions are included.

implicitly follow the rules, even though we have not directly accepted them. Our acceptance was *implicit*. This is similar to how laws in a society are tacitly accepted by being members of society. We have not *explicitly* accepted every law, but by engaging in society, we nonetheless conform to these laws.

Reiland (2023) mentions that certain activities might be constituted by some default rules. For example, Williamson (1996) says that "In natural languages, the default use of declarative sentences is to make assertions" (p. 511). Thus, when one makes an assertion, one is, by default, under the force of this rule, without one having to consciously accepting it. More exactly, "to continue accepting the rule, you do not have to do anything beyond not opting out" (Reiland 2023, p. 576).

Harry will implicitly follow the (UI) rule when learning the seeking-game, as the rule will automatically hold due to the setup of the game. As discussed in the previous section, he will further be in a position to recognise that (UI) holds in the game, as he can think in terms of winning strategies, and realise that the (UI) rule, spelled out as winning strategy preservation, will make sense for him. He will win in the particular instance *because* he can win in the universal case. There will be a graspable common cause for him which will allow him to form the concept All, and associate the (UI) rule to it. He will then be able to follow the (UI) rule if he so chooses, i.e., if he deems it useful in his situation.

5.3.2 *We need (UI) in order to learn the game*

A simple worry might arise. Won't Harry still need to perform (UI) in order to learn the rules of Hintikka's game? This is because, on most models, rules are seen as universal and conditional (Finn 2021). Thus, despite not giving rules that presuppose knowledge of the defined logical constant, it could be argued that an application of what is supposed to be learned nonetheless lurks in the background.

Firstly, I believe one source of concern is unfounded, namely the idea that since rules are conditional and universal, (UI) must be lurking in the background. This would be the case if Harry would be told the game as a set of formal rules. However, I believe these games can be learned simply through mimicking the behaviour of some supposed teacher, and thus learning to play is not a matter of learning some other rules, but of being sensitive to one's behaviour, and being able to mimic what others are doing. For example, I do not learn how to dance by being given a list with instructions which I then execute, but the instructor dances and I copy their moves. This way of learning bypasses the need for formal rules.[22]

Learning to infer might not have such a behavioural counterpart. Mimicking someone's act of inferring a conclusion from a set of premisses would

[22] Indeed, it can be seen in a rather know-how rather than a know-that fashion.

not count as an inference. That is because inference is taken to be more than a simple association between believing/asserting premisses and believing/asserting a conclusion, but we usually believe/assert the latter *because* we believe/assert the former. To use Boghossian's (2014, p. 4f) terminology, we *take* the premisses to support the conclusion. Mimicking someone's assertions of premisses and conclusions lacks this condition, or any sort of link between the two actions, and thus we cannot take someone imitating someone's inferential performance to also be performing an inference. Indeed, intentional aspects of action are a thorny point for a behavioural analysis, as Tomasello and Carpenter (2005) point out, since "humans understand the behavior of others not just as body movements, but as intentional, goal-directed action" (p. 133).

Yet don't we also need a similar condition when learning how to search for something? Mimicking someone's searching behaviour, could be said, is not an act of searching, but an act of mimicking. I don't think this is a very strong reply. For example, when we mimic someone's jumping, we indeed mimic, but don't we also jump as well? Nonetheless, I think a stronger objection can be made based on this, namely that usually searching has a *goal*, and thus one will be able to recognise when this goal has been achieved and stop the searching activity. In that sense, searching is a bit closer to inferring than to jumping, and one could wonder if figuring out that searching has a goal, and then figuring out what the goal might be, is an act of inference, and thus subject to the use of (UI) somewhere in the process.

Thus, for Harry to properly grasp the quantifier, not only does he need to understand what it is to search-for-something, but also to understand that searching is a goal-oriented behaviour, and the goal must be something inferred from observing others' actions. For simplicity, I will assume that Harry will be able to recognise whether a goal has been achieve, e.g., whether an object is indeed blue, or whether something is indeed a triangle. In other words, he can observe things just fine, and properly classify them.

Let's briefly take stock. Harry needs to learn two things in order to play the game. First, he needs to learn what searching behaviour looks like. In that sense, he needs to figure out how to search for something, and this, I believe, can be done by mere imitation.[23] Secondly, and more importantly, Harry needs to realise that the searching activity has a goal, and that goal is finding a way to verify the sentence. In other words, he has to connect the behaviour to the goal, not much different from how children learn new words:

> when a child hears a speaker use an unfamiliar term, the child must use his or her understanding of the communicative situation to focus his or her attention on the same entity (Akhtar and Tomasello 2000, p. 116).

[23] Or, if one is weary of imitation for whatever reasons, we can try and teach him how to search via examples, see Pettit (1990).

> sharing an understanding with their social partners about the overall goal of the activity [...] may allow infants to comprehend some of the language used within the activity (Akhtar and Tomasello 2000, p. 124).

Now, how might Harry learn the goal of the activity? First, he will have to recognise that someone is trying to teach him something, that someone's actions have a goal. However, this can be done by ostention, and this sort of process is non-inferential (Csibra 2010; Fiebich and Coltheart 2015). Then, how might Harry learn that *finding* is the goal of the activity of searching? A first natural instinct would be for the teacher to engage in the searching behaviour, and mark the end of the process. For example, say I want to check whether all things in this box are blue. I might start to pick up all things one by one, check them, put them back, and stop, say, when all the balls are exhausted, or I find a non-blue ball. In both cases, the teacher can mark the end of the activity in an explicit manner, say, by shouting 'Done'. I think explicitly marking the end of the activity is helpful, as such markers can aid in understanding what the activity is trying to achieve: "infants as young as 14 months of age can use verbal or emotional markers of intentionality to determine what an adult is trying to do [...] 18-month-olds respond differently when the action is accompanied by satisfied vs. dissatisfied emotional expressions and vocalizations" (Carpenter and Call 2007, p. 142, references omitted). I imagine it might be helpful to also mark a successful search vs. a non-successful one in a similar vein, perhaps with a 'Yes' to indicate the search marks all objects as blue, and 'No' to indicate the finding of a non-blue item. From this, Harry should come to understand that the verification of a quantified sentence consists in the activity of searching, and what a successful search looks like. The stopping of the activity when something is found, or when the space is exhausted, should help in identifying the goal of the activity (see also the experiments summarised in Tomasello and Carpenter 2005, §5.3.5).

However, one might not be convinced that the above process still does not use (UI) somewhere. After all, the literature says that the agent needs to *Infer the goal of the demonstrator* (Carpenter and Call 2007, p. 144). However, I do not think one needs to worry too much. First, the kind of inference at play is most likely not a deduction, but an induction (i.e., an inference to the best explanation). The learner tries to guess, based on salient information, the goal of the demonstrator. But this might seem like a shallow response, so let me strengthen it. In the case of learning (UI), we are going from a universal to a particular, i.e., from a more informational sentence to a less informational sentence. However, in the case of learning, one has to go the other way around. From limited examples, one has to generalize the rules, and the goal.[24] Since

[24]Therefore, the concern should be not whether one uses (UI) in order to learn the game,

multiple different activities might have the same goal, and the same goal might be achieved through wildly differing activities, the inference is not an inference in the deductive sense, but could better be captured by other terms such as 'abstraction', 'generalization', and so on.

Of course, what I have provided is not intended as being a knock-down argument against someone suggesting that (UI) is a must for learning the game, but only some points against suggesting that must be so. If I am right, then Harry can adopt the required *concept*, and automatically come to follow the rules of the concept. If I am more than right, this can be done through learning the aforementioned seeking games. If not, then I am confident that other learning methods can be conceived.

5.4 A SKETCH FORWARD

What I have provided for Harry is a sketch of how the (AP) can be resisted at the individual level. There are thus many questions that must be left, for the moment, unanswered, and some idealizations that must remain in place for this thesis to have a reasonable length. Importantly, since I take Harry to have managed to learn (a non-quantificational fragment of) English, I assume he is able to follow rules, and thus I set any rule-following considerations aside. The genealogical method, as mentioned, has already been applied to the issue of rule-following (Pettit 2023; Weiss 2022). If one is sceptical of Harry being able to follow the rule, or of being able to acquire the same concept as us, then this worry is not one that arises just for my proposal, but could underlie any proposal. For example, if Harry were to adopt, say, the concept QUALL (analogous to the quus in Kripke 1982) instead of ALL.

One issue on which I remained silent is the rational reconstruction itself. In other words, I have stopped well before even sketching a true reconstruction of quantification, which would require its own separate work. However, by shifting the focus to how Harry can learn the quantifiers, I believe the ground is set for continuing in this direction. Importantly, by pointing this direction, we will be in a better position to answer the question plaguing the proposal of Devitt and Roberts (2024), namely the logic coach regress.

It could be that the universal quantifier arises due to the need of predicating properties for unencountered objects.[25] For example, in saying *All humans are mortal* I not only aim to make a statement about current humans, but also about future humans, or humans that I have not met. Another possibility could be that it arises due to a need to express infinite conjunctions, or to signal in one fell swoop endorsement that could not be expressed in a timely

but rather whether one will come to follow a deviant rule, or form a deviant concept (Kripke 1982).

[25] A similar point is raised regarding the use of schemas (Fraser and Mount 2020), but I will ignore this, since the concept of a schema is too technical to appear naturally in spoken language.

manner. For example, it would be impractical to say *2 is divisible by 2*, *4 is divisible by 2*, … But, in order to express my commitment, I could utter *All even numbers are divisible by 2*.

I also want to dispel some potential worries with my approach. First of all, by employing such a method I am not committing a *genetic fallacy*, i.e., conflating genesis with justification. I am in no way putting forward the method in order to justify the way we currently use certain logical tools, in particular to justify one reading or another of the meaning of the quantifiers. I am simply asking why a community might use quantification, and showing how it might naturally arise. In this sense, I am leaving open forward-looking ameliorative disputes. As Kuusela points out, "[o]rigin does not fix future uses and logical status" (Kuusela 2019, p. 190). An important caveat to mention is that it might be doubtful that a society will evolve a quantifier such as the mathematical All outside of a particular cultural practice. That is, a quantifier aiming at universality understood as exception-less generality.

Of course, there could be much more to be said about the possible role of quantifiers, for example their interplay with numerical cognition (Szymanik et al. 2023), their emergence in actual languages (Kocab et al. 2022), or their cognitive aspects (Knowlton et al. 2022). A proper and rigorous reconstruction will need to take many such aspects into account when building up a convincing narrative. I believe this can be done, but will not the place to do it. I am nonetheless confident that the picture sketched in this chapter is enough as a proof-of-concept.

CHAPTER 6

Conclusions and takeaways

The Adoption Problem is a misleading problem. If viewed, as originally presented, as a problem at the level of rules, the reply is simple: Harry doesn't understand the rule. If presented as a problem at the level of meanings/concepts, it is straightforward. Presenting someone with a rule they have never seen before is not a reason to think that failure to adopt the rule means it must have a special epistemic status. The natural response, as I have argued in this thesis, is that the person simply failed to understand the rule. This was my diagnosis of Harry. In order to acquire the concepts used in the formulation of the rule, Harry must be immersed in some practice in which the rule *makes sense*. Without such practice, the rule itself will not make sense to him. The Adoption Problem arises from the misleading assumption that it will.

Birman could reply that my method goes beyond her assumptions. More exactly, her argument was more modest, namely that under the conditions demanded by adoption, basic rules are unadoptable. By going beyond those assumptions, I am attacking a straw-man. However, my argument does two things. Firstly, it shows how one could adopt a basic concept, and this indeed requires methods that go beyond the boundaries of Birman's setting. However, this is what I think makes the problem more interesting. As I have argued, Harry failed to adopt (UI) not because the rule is unadoptable, but because he failed to understand what the rule says. Then, Birman's problem is a false problem. No one learns rules in the way Harry is supposed to. Nonetheless, we could stay withing Birman's setting, but I do not think there's anything interesting to be found there.

The Adoption Problem arose, I believe, from the idea that someone could both understand a certain concept, but not be able to apply it in the most basic

of cases. In such cases, I believe, our intuition that the person failed to grasp the rule is the right diagnosis, and in Chapter 4 of my thesis I showed how we can substantiate this intuition by arguing that not endorsing certain constitutive inferences is a sign of misunderstanding. Contra Williamson (2022), someone who rejects these constitutive inferences, far from understanding the expression the same way we do, actually uses the expression with a different meaning. Someone who rejects that *Everything is self-identical* entitles one to infer *Frank is self-identical*, far from having a peculiar kind of understanding, has a misunderstanding of the way we use the concept All.

This misunderstanding, I claimed, can be solved by immersing Harry in a practice in which the quantifiers are embedded in a natural way, which was the positive contribution of my thesis in Chapter 5. This way of learning is not conditioning, but a way of teaching the quantifiers which gives them importance in a practical sense. Harry, by learning to play Hintikka's games, will be in a position to infer in accordance with, and accept, (UI). Further, this is not running counter to Birman's claim that an inferential practice must already be present. What is needed is a practice which will give rise to the concept All, but this practice need not explicitly have the (UI) rule as an essential statement, but is simply a consequence of our use of quantifiers. This is similar to how we first learn a language, and then follow the rules of grammar, or learn to count, and then automatically infer in accordance with Hume's Principle. We are not explicitly told (UI), nor the rules of grammar, nor Hume's Principle, but in virtue of the activities we immerse ourselves in, we will automatically come to conform to these principles and rules, and be in a position to recognise that the principles hold in our practices.

On my view, is 'adopting a logic' still an incoherent notion as Kripke would claim? Not quite. I believe we can make sense of changes in logic by thinking of these in terms of conceptual changes. The reason certain inferences or rules might seem 'intuitive' is due to their role in certain activities, in the case of quantifiers those of 'seeking-and-finding'. Then it is thoroughly unsurprising that we find the (UI) rule intuitive *once we learn the game of 'seeking-and-finding'*, similarly to how we find Hume's Principle compelling once we learn how to count. The fact that we find certain rules or principles primitively compelling is a mixture of being immersed in certain activities which conform to those rules, and our psychological build-up. If someone has never seen a word or does not possess a concept used to phrase some rule, they will be in no position to find anything related to that word/concept 'intuitive'.

Finally, one potential lesson of my argument might be a thoroughly pragmatic one: it's not the case that certain principles have a special epistemic status, but that certain activities have a special status for us, and then in turn certain principles are constitutive of those certain activities only in virtue of how we set up those practices. These practices help one understand and acquire the concepts that then obey certain rules, but importantly one will not be able to acquire basic concepts by being told rules for using that concept,

similarly to how someone who has never seen the colour green will not be able to point out a green object. Of course, which activities might be basic in this sense is an open question, one which I leave for future, likely anthropological works, to answer.

Bibliography

Akhtar, Nameera and Michael Tomasello (2000), "The Social Nature of Words and Word Learning", in *Becoming a Word Learner: A Debate on Lexical Acquisition*, ed. by Roberta Micknick Golinkoff and Kathryn Hirsh-Pasek, Oxford University Press, pp. 115-135. (Cit. on pp. 83, 84.)

Arnold, Jack and Stewart Shapiro (2007), "Where in the (World Wide) Web of Belief is the Law of Non-Contradiction?", *Noûs*, vol. 41, no. 2, pp. 276-297, DOI: 10.1111/j.1468-0068.2007.00647.x. (Cit. on p. 19.)

Asay, Jamin (2021), "Something is True", *Philosophy and Phenomenological Research*, vol. 105, no. 3, pp. 687-705, DOI: 10.1111/phpr.12836. (Cit. on p. 46.)

Balcerak Jackson, Magdalena and Brendan Balcerak Jackson (2012), "Understanding and Philosophical Methodology", *Philosophical Studies*, vol. 161, no. 2, pp. 185-205, DOI: 10.1007/s11098-011-9729-y. (Cit. on p. 44.)

Barker-Plummer, David, Jon Barwise, and John Etchemendy (2007), *Tarski's World: Revised and Expanded*, Center for the Study of Language and Information - Lecture Notes, CSLI Publications. (Cit. on p. 70.)

Başkent, Can (2016), "Game Theoretical Semantics for Some Non-Classical Logics", *Journal of Applied Non-Classical Logics*, vol. 26, no. 3, pp. 208-239, DOI: 10.1080/11663081.2016.1225488. (Cit. on p. 72.)

Bealer, George (1998), "A Theory of Concepts and Concepts Possession", *Philosophical Issues*, vol. 9, pp. 261-301, DOI: 10.2307/1522977. (Cit. on p. 51.)

Beck, Jacob (2015), "Analogue Magnitude Representations: A Philosophical Introduction", *British Journal for the Philosophy of Science*, vol. 66, no. 4, pp. 829-855, DOI: 10.1093/bjps/axu014. (Cit. on p. 64.)

Bengson, John and Marc Moffett (2007), "Know-How and Concept Possession", *Philosophical Studies*, vol. 136, no. 1, pp. 31-57, DOI: 10.1007/s11098-007-9146-4. (Cit. on p. 56.)

Berger, Alan (2011), "Kripke on the Incoherency of Adopting a Logic", in, *Saul Kripke*, ed. by Alan Berger, Cambridge University Press, pp. 177-208. (Cit. on p. 3.)

Berry, Sharon (2013), "Default Reasonableness and the Mathoids", *Synthese*, vol. 190, no. 17, pp. 3695-3713, DOI: 10.1007/s11229-012-0219-3. (Cit. on pp. 39, 71.)

Berto, Francesco (2015), "A Modality Called 'Negation'", *Mind*, vol. 124, no. 495, pp. 761-793, DOI: 10.1093/mind/fzv026. (Cit. on p. 66.)

Berto, Francesco and Greg Restall (2019), "Negation on the Australian Plan", *Journal of Philosophical Logic*, vol. 48, no. 6, pp. 1119-1144, DOI: 10.1007/s10992-019-09510-2. (Cit. on p. 66.)

Besson, Corine (2018), "Norms, Reasons and Reasoning: A Guide Through Lewis Carroll's Regress Argument", in *The Oxford Handbook of Reasons and Normativity*, ed. by Daniel Star, Oxford University Press, pp. 504-528. (Cit. on p. 13.)

— (2019), "Knowledge of Logical Generality and the Possibility of Deductive Reasoning", in *Inference and Consciousness*, ed. by Anders Nes and Timothy Hoo Wai Chan, Routledge, pp. 172-196. (Cit. on pp. 29, 31, 78, 79.)

— (forthcoming), *Logic, Reasoning and Carroll's Regress*, Oxford University Press. (Cit. on p. 29.)

Besson, Corine and Anandi Hattiangadi (m.s.), "Deductive Reasoning Without Rule-following". (Cit. on p. 78.)

Birman, Romina (2024), "The Adoption Problem and the Epistemology of Logic", *Mind*, vol. 133, no. 529, pp. 37-60, DOI: 10.1093/mind/fzad009. (Cit. on pp. 2-4, 9, 12, 16, 21, 28, 30, 34, 69, 80.)

Blake-Turner, Christopher (2022), "The Hereby-Commit Account of Inference", *Australasian Journal of Philosophy*, vol. 100, no. 1, pp. 86-101, DOI: 10.1080/00048402.2020.1843062. (Cit. on p. 23.)

Boghossian, Paul (2000), "Knowledge of Logic", in *New Essays on the A Priori*, ed. by Paul Boghossian and Christopher Peacocke, Oxford University Press, pp. 229-254. (Cit. on pp. 8, 16.)

— (2003), "Blind Reasoning", *Supplement to the Proceedings of the Aristotelian Society*, vol. 77, no. 1, pp. 225-248, DOI: 10.1111/1467-8349.00110. (Cit. on pp. 16, 33.)

— (2012), "Inferentialism and the Epistemology of Logic: Reflections on Casalegno and Williamson", *Dialectica*, vol. 66, no. 2, pp. 221-236, DOI: 10.1111/j.1746-8361.2012.01303.x. (Cit. on p. 53.)

— (2014), "What is Inference?", *Philosophical Studies*, vol. 169, no. 1, pp. 1-18, DOI: 10.1007/s11098-012-9903-x. (Cit. on pp. 23, 83.)

Boghossian, Paul and Crispin Wright (2024), "Kripke, Quine, the 'Adoption Problem' and the Empirical Conception of Logic", *Mind*, vol. 133, no. 529, pp. 86-116, DOI: 10.1093/mind/fzad011. (Cit. on pp. 35, 57, 78.)

Brandom, Robert (1993), "The Social Anatomy of Inference", *Philosophy and Phenomenological Research*, vol. 53, no. 3, pp. 661-666, DOI: 10.2307/2108089. (Cit. on p. 37.)

— (1994), *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Harvard University Press. (Cit. on p. 53.)
— (2000), *Articulating Reasons: An Introduction to Inferentialism*, Cambridge: Harvard University Press. (Cit. on pp. 52, 53.)
Brîncuş, Constantin C. (2024), "Inferential Quantification and the ω-Rule", in *Perspectives on Deduction: Contemporary Studies in the Philosophy, History and Formal Theories of Deduction*, ed. by Antonio Piccolomini D'Aragona, Springer Verlag, pp. 345-372. (Cit. on p. 39.)
Brun, Georg (2022), "Re-Engineering Contested Concepts. A Reflective-Equilibrium Approach", *Synthese*, vol. 200, no. 2, pp. 1-29, DOI: 10.1007/s11229-022-03556-7. (Cit. on pp. 43, 47, 48.)
Buijsman, Stefan (2017), "Learning the Natural Numbers as a Child", *Noûs*, vol. 53, no. 1, pp. 3-22, DOI: 10.1111/nous.12219. (Cit. on pp. 11, 60, 64.)
Burge, Tyler (1979), "Individualism and the Mental", *Midwest Studies in Philosophy*, vol. 4, no. 1, pp. 73-122, DOI: 10.1111/j.1475-4975.1979.tb00374.x. (Cit. on pp. 44, 49, 57.)
Burgess, Alexis, Herman Cappelen, and David Plunkett (eds.) (2020), *Conceptual Engineering and Conceptual Ethics*, New York, USA: Oxford University Press. (Cit. on p. 48.)
Cappelen, Herman (2018), *Fixing Language: An Essay on Conceptual Engineering*, Oxford: Oxford University Press. (Cit. on p. 48.)
Carey, Susan (2009a), *The Origin of Concepts*, Oxford University Press. (Cit. on pp. 11, 64, 65.)
— (2009b), "Where Our Number Concepts Come From", *Journal of Philosophy*, vol. 106, no. 4, pp. 220-254, DOI: 10.5840/jphil2009106418. (Cit. on p. 64.)
Carnap, Rudolf (1937), *The Logical Syntax of Language*, ed. by Amethe Smeaton, London, K. Paul, Trench, Trubner & co. (Cit. on p. 16.)
— (1950), *Logical Foundations of Probability*, The University of Chicago Press. (Cit. on p. 43.)
Carpenter, Malinda and Josep Call (2007), "The question of 'what to imitate': inferring goals and intentions from demonstrations", in, *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*, ed. by Chrystopher L. Nehaniv and Kerstin Dautenhahn, Cambridge University Press, pp. 135-152. (Cit. on p. 84.)
Carroll, Lewis (1895), "What the Tortoise Said to Achilles", *Mind*, vol. 4, no. 14, pp. 278-280, DOI: 10.1093/mind/iv.14.278. (Cit. on pp. 4, 12, 13.)
Casalegno, Paolo (2004), "Logical Concepts and Logical Inferences", *Dialectica*, vol. 58, no. 3, pp. 395-411, DOI: 10.1111/j.1746-8361.2004.tb00313.x. (Cit. on p. 52.)
Chalmers, David J. (2011), "Verbal Disputes", *Philosophical Review*, vol. 120, no. 4, pp. 515-566, DOI: 10.1215/00318108-1334478. (Cit. on p. 46.)

Chalmers, David J. and Frank Jackson (2001), "Conceptual Analysis and Reductive Explanation", *Philosophical Review*, vol. 110, no. 3, pp. 315-61, DOI: 10.1215/00318108-110-3-315. (Cit. on p. 44.)

Cherniak, Christopher (1986), *Minimal Rationality*, MIT Press. (Cit. on pp. 10, 37, 40, 71.)

Chudnoff, Elijah (2024), "Inferential Seemings", in *Oxford Studies in Philosophy of Mind: Volume 4*, ed. by Uriah Kriegel, Oxford University Press, pp. 311-333, ISBN: 9780198924128. (Cit. on p. 24.)

— (forthcoming), "Reasoned Change in Logic", in *Evidentialism at 40: New Arguments, New Angles*, ed. by Scott Stapleford, Kevin McCain, and Matthias Steup, Routledge. (Cit. on pp. 23, 25, 31.)

Cook Wilson, John (1926), *Statement and Inference: With Other Philosophical Papers*, ed. by A. S. L. Farquharson, Oxford University Press. (Cit. on p. 13.)

Crimmins, Mark (1989), "Having Ideas and Having the Concept", *Mind and Language*, vol. 4, no. 4, pp. 280-294, DOI: 10.1111/j.1468-0017.1989.tb00257.x. (Cit. on p. 51.)

Csibra, Gergely (2010), "Recognizing Communicative Intentions in Infancy", *Mind and Language*, vol. 25, no. 2, pp. 141-168, DOI: 10.1111/j.1468-0017.2009.01384.x. (Cit. on p. 84.)

Davidson, Donald (1968), "On Saying That", *Synthese*, vol. 19, no. 1-2, pp. 130-146, DOI: 10.1007/bf00568054. (Cit. on p. 44.)

De Cruz, Helen (2018), "Testimony and Children's Acquisition of Number Concepts", in *Naturalizing Logico-Mathematical Knowledge*, ed. by Sorin Bangu, Routledge, pp. 164-178. (Cit. on p. 65.)

Dehaene, Stanislas (1997), *The Number Sense : How the Mind Creates Mathematics: How the Mind Creates Mathematics*, Oxford University Press. (Cit. on p. 64.)

Devitt, Michael (2006), *Ignorance of Language*, Oxford University Press. (Cit. on p. 25.)

Devitt, Michael and Jillian Rose Roberts (2024), "Changing Our Logic: A Quinean Perspective", *Mind*, vol. 133, no. 529, pp. 61-85, DOI: 10.1093/mind/fzad010. (Cit. on pp. 10, 21-23, 30, 31, 75, 79, 80, 85.)

Dogramaci, Sinan (2010), "Knowledge of Validity", *Noûs*, vol. 44, no. 3, pp. 403-432, DOI: 10.1111/j.1468-0068.2010.00746.x. (Cit. on p. 8.)

Drobňák, Matej (2024), "Communication Without Shared Meanings", *Acta Analytica*, pp. 1-18, DOI: 10.1007/s12136-024-00604-4. (Cit. on p. 37.)

Dummett, Michael (1973), *Frege: Philosophy of Language*, London: Duckworth. (Cit. on pp. 18, 72.)

— (1978), "The Justification of Deduction", in *Truth and other enigmas*, Harvard University Press. (Cit. on p. 8.)

— (1991), *The Logical Basis of Metaphysics*, Harvard University Press. (Cit. on p. 38.)

Edwards, Jim (1995), "The Universal Quantifier and Dummett's Verificationist Theory of Sense", *Analysis*, vol. 55, no. 2, pp. 90-97, DOI: 10.1093/analys/55.2.90. (Cit. on p. 38.)

Enoch, David and Joshua Schechter (2008), "How Are Basic Belief-Forming Methods Justified?", *Philosophy and Phenomenological Research*, vol. 76, no. 3, pp. 547-579, DOI: 10.1111/j.1933-1592.2008.00157.x. (Cit. on p. 39.)

Evnine, Simon J. (2008), *Epistemic Dimensions of Personhood*, Oxford University Press. (Cit. on pp. 40, 70, 75.)

Fiebich, Anika and Max Coltheart (2015), "Various Ways to Understand Other Minds: Towards a Pluralistic Approach to the Explanation of Social Understanding", *Mind and Language*, vol. 30, no. 3, pp. 235-258, DOI: 10.1111/mila.12079. (Cit. on p. 84.)

Field, Hartry (1998), "Epistemological Nonfactualism and the a Prioricity of Logic", *Philosophical Studies*, vol. 92, no. 1, pp. 1-24, DOI: 10.1023/a:1017147216094. (Cit. on p. 26.)

Finley, Robby (2024), "Towards a pragmatist epistemology for theory choice in logic", *Synthese*, vol. 204, no. 1, pp. 1-28, DOI: 10.1007/s11229-024-04654-4. (Cit. on pp. 26, 31, 67.)

Finn, Suki (2019), "The Adoption Problem and Anti-Exceptionalism About Logic", *Australasian Journal of Logic*, vol. 16, no. 7, pp. 231-249, DOI: 10.26686/ajl.v16i7.5916. (Cit. on p. 2.)

— (2021), "Limiting Logical Pluralism", *Synthese*, vol. 198, no. Suppl 20, pp. 4905-4923, DOI: 10.1007/s11229-019-02134-8. (Cit. on pp. 2, 82.)

Fodor, Jerry and Ernie Lepore (2007), "Brandom Beleaguered", *Philosophy and Phenomenological Research*, vol. 74, no. 3, pp. 677-691, DOI: 10.1111/j.1933-1592.2007.00045.x. (Cit. on p. 69.)

Fodor, Jerry A. (1975), *The Language of Thought*, Harvard University Press. (Cit. on p. 34.)

— (1981), "The Present Status of the Innateness Controversy", in *Representations: Philosophical essays on the foundations of cognitive science*, ed. by Jerry A. Fodor, MIT Press, pp. 257-316. (Cit. on p. 34.)

— (2008), *LOT 2: The Language of Thought Revisited*, New York: Oxford University Press. (Cit. on p. 34.)

Frápolli, María José (2023), *The Priority of Propositions. A Pragmatist Philosophy of Logic*, Springer Verlag. (Cit. on p. 37.)

Fraser, Rachel Elizabeth and Beau Madison Mount (2020), "Absolutely General Knowledge", *Philosophy and Phenomenological Research*, vol. 103, no. 3, pp. 547-566, DOI: 10.1111/phpr.12710. (Cit. on p. 85.)

Gallie, W. B. (1956), "Essentially Contested Concepts", *Proceedings of the Aristotelian Society*, vol. 56, no. 1, pp. 167-198, DOI: 10.1093/aristotelian/56.1.167. (Cit. on p. 43.)

Gamester, Will (2023), "Nothing is True", *Journal of Philosophy*, vol. 120, no. 6, pp. 314-338, DOI: 10.5840/jphil2023120612. (Cit. on p. 46.)

Gauker, Christopher (1997), "Universal Instantiation: A Study of the Role of Context in Logic", *Erkenntnis*, vol. 46, no. 2, pp. 185-214, DOI: 10.1023/a:1005380919258. (Cit. on p. 41.)

Girard, Jean-Yves (1999), "On the Meaning of Logical Rules I: Syntax Versus Semantics", in *Computational Logic*, ed. by Ulrich Berger and Helmut Schwichtenberg, Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 215-272, ISBN: 978-3-642-58622-4. (Cit. on p. 69.)

Goldfarb, Warren D. (1979), "Logic in the Twenties: The Nature of the Quantifier", *Journal of Symbolic Logic*, vol. 44, no. 3, pp. 351-368, DOI: 10.2307/2273128. (Cit. on p. 16.)

Goodman, Jeremy (2024), "Higher-Order Logic as Metaphysics", in *Higher-Order Metaphysics*, ed. by Peter Fritz and Nicholas K. Jones, Oxford University Press, pp. 73-106. (Cit. on p. 38.)

Grice, H. P. and P. F. Strawson (1956), "In Defense of a Dogma", *Philosophical Review*, vol. 65, no. 2, pp. 141-158, DOI: 10.2307/2182828. (Cit. on pp. 38, 44.)

Hale, Bob (2002), "Basic Logical Knowledge", *Royal Institute of Philosophy Supplement*, vol. 51, pp. 279-304, DOI: 10.1017/s1358246100008171. (Cit. on p. 37.)

Hale, Bob and Crispin Wright (2001), *The Reason's Proper Study: Essays Towards a Neo-Fregean Philosophy of Mathematics*, Clarendon Press. (Cit. on pp. 11, 59.)

Hand, Michael (1988), "How Game-Theoretical Semantics Works: Classical First-Order Logic", *Erkenntnis*, vol. 29, no. 1, pp. 77-93, DOI: 10.1007/bf00166366. (Cit. on p. 75.)

Hanna, Robert (2006), *Rationality and Logic*, The MIT Press. (Cit. on pp. 13, 16.)

Hannon, Michael (2021), "Recent Work in the Epistemology of Understanding", *American Philosophical Quarterly*, vol. 58, no. 3, pp. 269-290, DOI: 10.2307/48616060. (Cit. on p. 36.)

Harman, Gilbert (1986), *Change in View: Principles of Reasoning*, Cambridge, MA, USA: MIT Press. (Cit. on p. 27.)

Hart, H. L. A. (1961), *The Concept of Law*, Oxford, United Kingdom: Oxford University Press. (Cit. on pp. 37, 61.)

Hartonas, Chrysafis (2019), "Game-Theoretic Semantics for Non-Distributive Logics", *Logic Journal of the IGPL*, vol. 27, no. 5, pp. 718-742, DOI: 10.1093/jigpal/jzy079. (Cit. on p. 72.)

Hattiangadi, Anandi (2023), "Logical Conventionalism and the Adoption Problem", *Aristotelian Society Supplementary Volume*, vol. 97, no. 1, pp. 47-81, DOI: 10.1093/arisup/akad008. (Cit. on p. 22.)

Heck, Richard G. (2011), *Frege's Theorem*, New York: Clarendon Press. (Cit. on p. 59.)

Hintikka, Jaakko (1973), *Logic, Language-Games and Information: Kantian Themes in the Philosophy of Logic*, Oxford: Clarendon Press. (Cit. on pp. 2, 67, 68, 70, 73.)

— (1983), "Semantical Games and Transcendental Arguments", in, *The Game of Language: Studies in Game-Theoretical Semantics and Its Applications*, Dordrecht: Springer Netherlands, pp. 33-46, DOI: 10.1007/978-94-010-9847-2_2. (Cit. on p. 75.)

Hintikka, Jaakko and Gabriel Sandu (1997), "Game-Theoretic Semantics", in *Handbook of Logic and Language*, ed. by Johan van Benthem and Alice ter Meulen, Elsevier, pp. 361-410. (Cit. on p. 68.)

Hjortland, Ole Thomassen (2017), "Anti-Exceptionalism About Logic", *Philosophical Studies*, vol. 174, no. 3, pp. 631-658, DOI: 10.1007/s11098-016-0701-8. (Cit. on pp. 17, 20.)

Hlobil, Ulf (2019), "Inferring by Attaching Force", *Australasian Journal of Philosophy*, vol. 97, no. 4, pp. 701-714, DOI: 10.1080/00048402.2018.1564060. (Cit. on p. 23.)

— (2020), "Limits of Abductivism About Logic", *Philosophy and Phenomenological Research*, vol. 103, no. 2, pp. 320-340, DOI: 10.1111/phpr.12707. (Cit. on p. 15.)

Hodges, Wilfrid (2001), "Dialogue Foundations: A Sceptical Look", *Aristotelian Society Supplementary Volume*, vol. 75, no. 1, pp. 17-32, DOI: 10.1111/1467-8349.00076. (Cit. on pp. 70, 73.)

Horsten, Leon and Volker Halbach (2015), "Norms for Theories of Reflexive Truth", in *Unifying the Philosophy of Truth*, ed. by T. Achourioti, H. Galinon, J. Martínez Fernández, and K. Fujimoto, Springer, pp. 263-280. (Cit. on p. 46.)

Incurvati, Luca (2020), *Conceptions of Set and the Foundations of Mathematics*, Cambridge University Press. (Cit. on p. 37.)

— (forthcoming), "Three Kinds of Logical Expressivism", in *Oxford Handbook of the Philosophy of Logic*, ed. by F. Ferrari, E. Brendel, M. Carrara, O. Hjortland, G. Sagi, G. Sher, and F. Steinberger, Oxford University Press. (Cit. on p. 13.)

Incurvati, Luca and Julian J. Schlöder (2023), *Reasoning with Attitude: Foundations and Applications of Inferential Expressivism*, Oxford University Press. (Cit. on pp. 36, 37.)

Jenkins, C. S. I. (2014), "Merely Verbal Disputes", *Erkenntnis*, vol. 79, no. 1, pp. 11-30, DOI: 10.1007/s10670-013-9443-6. (Cit. on p. 46.)

Jenkins, Carrie (2008), *Grounding Concepts: An Empirical Basis for Arithmetical Knowledge*, Oxford University Press. (Cit. on p. 66.)

Knowlton, Tyler, Paul Pietroski, Justin Halberda, and Jeffrey Lidz (2022), "The Mental Representation of Universal Quantifiers", *Linguistics and Philosophy*, vol. 45, no. 4, pp. 911-941, DOI: 10.1007/s10988-021-09337-8. (Cit. on p. 86.)

Kocab, Annemarie, Kathryn Davidson, and Jesse Snedeker (2022), "The Emergence of Natural Language Quantification", *Cognitive Science*, vol. 46, no. 2, e13097, DOI: 10.1111/cogs.13097. (Cit. on p. 86.)

Kripke, Saul A. (1982), *Wittgenstein on Rules and Private Language: An Elementary Exposition*, Cambridge: Harvard University Press. (Cit. on pp. 6, 85.)

— (2024), "The Question of Logic", *Mind*, vol. 133, no. 529, pp. 1-36, DOI: 10.1093/mind/fzad008. (Cit. on pp. 2, 3, 5, 14, 17, 20, 21, 26, 28, 60.)

Kuhn, Thomas S. (1982), "Commensurability, Comparability, Communicability", *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, vol. 1982, pp. 669-688. (Cit. on p. 80.)

Kuusela, Oskari (2019), *Wittgenstein on Logic as the Method of Philosophy: Re-Examining the Roots and Development of Analytic Philosophy*, Oxford, England: Oxford University Press. (Cit. on p. 86.)

Laurence, Stephen and Eric Margolis (1999), "Concepts and Cognitive Science", in *Concepts: Core Readings*, ed. by Eric Margolis and Stephen Laurence, MIT Press, pp. 3-81. (Cit. on p. 40.)

Leech, Jessica (2015), "Logic and the Laws of Thought", *Philosophers' Imprint*, vol. 15, no. 12, pp. 1-27, http://hdl.handle.net/2027/spo.3521354.0015.012. (Cit. on pp. 37, 39.)

Liggins, David (2019), "In Defence of Radical Restrictionism", *Philosophy and Phenomenological Research*, vol. 98, no. 1, pp. 3-25, DOI: 10.1111/phpr.12391. (Cit. on p. 46.)

Linnebo, Øystein (2009), "The Individuation of the Natural Numbers", in *New Waves in Philosophy of Mathematics*, ed. by Ø. Linnebo O. Bueno, Palgrave-Macmillan, pp. 220-238. (Cit. on p. 59.)

Majer, Ondrej and Christian Fermüller (2018), "On Semantic Games for Łukasiewicz Logic", in *Jaakko Hintikka on Knowledge and Game Theoretical Semantics*, ed. by Hans van Ditmarsch and Gabriel Sandu, Springer, pp. 263-278. (Cit. on p. 72.)

Marcus, Eric (2020), "Inference as Consciousness of Necessity", *Analytic Philosophy*, vol. 61, no. 4, pp. 304-322, DOI: 10.1111/phib.12153. (Cit. on p. 23.)

Margolis, Eric (2020), "The Small Number System", *Philosophy of Science*, vol. 87, no. 1, pp. 113-134, DOI: 10.1086/706087. (Cit. on p. 65.)

Margolis, Eric and Stephen Laurence (2011), "Learning Matters: The Role of Learning in Concept Acquisition", *Mind and Language*, vol. 26, no. 5, pp. 507-539, DOI: 10.1111/j.1468-0017.2011.01429.x. (Cit. on p. 79.)

Martin, Ben (2021), "Identifying Logical Evidence", *Synthese*, vol. 198, no. 10, pp. 9069-9095, DOI: 10.1007/s11229-020-02618-y. (Cit. on p. 15.)

McDowell, John (1994), *Mind and World*, Cambridge: Harvard University Press. (Cit. on p. 53.)

McGee, Vann (1985), "A Counterexample to Modus Ponens", *Journal of Philosophy*, vol. 82, no. 9, pp. 462-471. (Cit. on pp. 42, 45.)
Miguens, Sofia (ed.) (2020), *The Logical Alien: Conant and His Critics*, Harvard University Press. (Cit. on p. 18.)
Miller, James (2022), "Merely Verbal Disputes and Common Ground", *Theoria*, no. 1, pp. 114-123, DOI: 10.1111/theo.12449. (Cit. on pp. 40, 46.)
Murzi, Julien and Florian Steinberger (2017), "Inferentialism", in *A Companion to the Philosophy of Language*, ed. by Bob Hale, Crispin Wright, and Alexander Miller, 2nd edition, Wiley-Blackwell, pp. 197-224. (Cit. on p. 35.)
Nolt, John (2021), "Free Logic", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Fall 2021, Metaphysics Research Lab, Stanford University. (Cit. on p. 41.)
Nutting, Eileen S. (2018), "The Limits of Reconstructive Neologicist Epistemology", *Philosophical Quarterly*, vol. 68, no. 273, pp. 717-738, DOI: 10.1093/pq/pqy010. (Cit. on p. 59.)
Padro, Romina (2015), *What the Tortoise Said to Kripke: the Adoption Problem and the Epistemology of Logic*, PhD thesis, CUNY Graduate Center, https://academicworks.cuny.edu/gc_etds/603/. (Cit. on pp. 2-7, 15, 20, 21, 28, 34.)
Paulson, Spencer (2024), "Internalizing Rules", *Philosophy and Phenomenological Research*, vol. 109, no. 2, pp. 630-649, DOI: 10.1111/phpr.13065. (Cit. on p. 74.)
Pavese, Carlotta (2021a), "Practical Concepts and Productive Reasoning", *Synthese*, vol. 199, no. 3-4, pp. 7659-7688, DOI: 10.1007/s11229-021-03132-5. (Cit. on p. 68.)
— (2021b), "Reasoning and Presuppositions", *Philosophical Topics*, vol. 49, no. 2, pp. 203-224, DOI: 10.5840/philtopics202149222. (Cit. on p. 19.)
— (2022), "Lewis Carroll's Regress and the Presuppositional Structure of Arguments", *Linguistics and Philosophy*, vol. 45, no. 1, pp. 1-38, DOI: 10.1007/s10988-020-09320-9. (Cit. on pp. 12, 13.)
Peacocke, Christopher (2024), "Deductive Inference and Mental Agency", *Analytic Philosophy*, pp. 1-13, DOI: 10.1111/phib.12360. (Cit. on p. 23.)
Peirce, Charles S. (1931), *Collected Papers Vols. 1-6*, ed. by Charles Hartshorne and Paul Weiss, Cambridge, Belknap Press of Harvard University Press. (Cit. on p. 27.)
Pettit, Philip (1990), "The Reality of Rule-Following", *Mind*, vol. 99, no. 393, pp. 1-21, DOI: 10.1093/mind/xcix.393.1. (Cit. on p. 83.)
— (2023), "A Pragmatic Genealogy of Rule-Following", in *Neopragmatism: Interventions in First-order Philosophy*, ed. by Joshua Gert, Oxford University Press, pp. 141-169. (Cit. on pp. 73, 85.)
Picollo, Lavinia and Thomas Schindler (2018), "Deflationism and the Function of Truth", *Philosophical Perspectives*, vol. 32, no. 1, pp. 326-351, DOI: 10.1111/phpe.12113. (Cit. on p. 48.)

Pietarinen, Ahti-Viekko (2006), *Signs of Logic: Peircean Themes on the Philosophy of Language, Games, and Communication*, Dordrecht, Netherland: Springer. (Cit. on p. 70.)
Price, Huw (1990), "Why 'Not?'", *Mind*, vol. 99, no. 394, pp. 221-238, DOI: 10.1093/mind/xcix.394.221. (Cit. on p. 72.)
— (2011), *Naturalism Without Mirrors*, Oxford University Press. (Cit. on p. 63.)
Priest, Graham (1979), "Two Dogmas of Quineanism", *Philosophical Quarterly*, vol. 29, no. 117, pp. 289-301, DOI: 10.2307/2219445. (Cit. on p. 13.)
— (2006), *Doubt Truth to Be a Liar*, Oxford University Press. (Cit. on p. 3.)
— (2014), "Revising logic", in, *The Metaphysics of Logic*, ed. by PenelopeEditor Rush, Cambridge University Press, pp. 211-223. (Cit. on p. 27.)
— (2016), "Logical Disputes and the *A Priori*", *Logique et Analyse*, no. 236, pp. 347-366, DOI: 10.2143/LEA.236.0.3186061. (Cit. on p. 16.)
Prior, Arthur (1960), "The Runabout Inference-Ticket", *Analysis*, vol. 21, no. 2, pp. 38-39, DOI: 10.1093/analys/21.2.38. (Cit. on p. 55.)
Putnam, Hilary (1968), "Is Logic Empirical?", *Boston Studies in the Philosophy of Science*, vol. 5. (Cit. on p. 3.)
— (1978), "There is at Least One a Priori Truth", *Erkenntnis*, vol. 13, no. 1, pp. 153-170, DOI: 10.1007/bf00160892. (Cit. on p. 44.)
Queloz, Matthieu (2019), "The Points of Concepts: Their Types, Tensions, and Connections", *Canadian Journal of Philosophy*, vol. 49, no. 8, pp. 1122-1145, DOI: 10.1080/00455091.2019.1584940. (Cit. on p. 72.)
Quine, Willard V. (1936), "Truth by Convention", in *Philosophical Essays for Alfred North Whitehead*, New York: Longman, Green and Co. (Cit. on p. 16.)
— (1951), "Two Dogmas of Empiricism", *Philosophical Review*, vol. 60, no. 1, pp. 20-43, DOI: 10.2307/2266637. (Cit. on pp. 3, 16, 18, 20.)
— (1986), *Philosophy of Logic (2nd Edition)*, Cambridge, MA: Harvard University Press. (Cit. on pp. 19, 44, 48.)
Rabin, Gabriel Oak (2020), "Toward a Theory of Concept Mastery: The Recognition View", *Erkenntnis*, vol. 85, no. 3, pp. 627-648, DOI: 10.1007/s10670-018-0040-6. (Cit. on pp. 38, 44.)
Reiland, Indrek (2023), "Rules of Use", *Mind and Language*, vol. 38, no. 2, pp. 566-583, DOI: 10.1111/mila.12404. (Cit. on p. 82.)
— (2024), "Rule-Following I: The Basic Issues (1st edition)", *Philosophy Compass*, vol. 19, no. 1, e12900, DOI: 10.1111/phc3.12900. (Cit. on p. 6.)
Resnik, Michael D. (2004), "Revising Logic", in *The Law of Non-Contradiction : New Philosophical Essays*, ed. by Graham Priest, JC Beall, and Bradley Armour-Garb, Oxford University Press, pp. 178-195. (Cit. on p. 27.)
Rips, Lance J., Amber Bloomfield, and Jennifer Asmuth (2008), "From Numerical Concepts to Concepts of Number", *Behavioral and Brain Sciences*, vol. 31, no. 6, pp. 623-642, DOI: 10.1017/s0140525x08005566. (Cit. on p. 64.)

Rumfitt, Ian (2015), *The Boundary Stones of Thought: An Essay in the Philosophy of Logic*, Oxford University Press. (Cit. on p. 41.)
Russell, Bertrand (2010), *Principles of Mathematics*, Routledge; trans. 1903. (Cit. on p. 13.)
Russell, Gillian (2015), "The Justification of the Basic Laws of Logic", *Journal of Philosophical Logic*, vol. 44, no. 6, pp. 793-803, DOI: 10.1007/s10992-015-9360-z. (Cit. on p. 16.)
Sabo, Dylan (2014), "Where Concepts Come From: Learning Concepts by Description and by Demonstration", *Erkenntnis*, vol. 79, no. 3, pp. 531-549, DOI: 10.1007/s10670-013-9522-8. (Cit. on p. 79.)
Samuels, Richard and Eric Snyder (2024), *Number Concepts: An Interdisciplinary Inquiry*, Cambridge University Press. (Cit. on pp. 64, 65.)
Sarnecka, Barbara W. (2021), "Learning to Represent Exact Numbers", *Synthese*, vol. 198, no. Suppl 5, pp. 1001-1018, DOI: 10.1007/s11229-015-0854-6. (Cit. on p. 65.)
Scharp, Kevin (2013), *Replacing Truth*, Oxford University Press. (Cit. on pp. 46-48.)
Searle, John (1980), "Minds, Brains, and Programs", *Behavioral and Brain Sciences*, vol. 3, no. 3, pp. 417-57, DOI: 10.1017/s0140525x00005756. (Cit. on p. 22.)
Shapiro, Stewart (2000), "The Status of Logic", in *New Essays on the A Priori*, ed. by Paul Boghossian and Christopher Peacocke, Oxford University Press, pp. 333-366. (Cit. on pp. 16, 20.)
Shaw, James R. (m.s.), "Inference and the Foundations of logic", Draft of 7 December 2023, https://www.jshaw.net/documents/IFL.pdf. (Cit. on p. 29.)
Sheffer, Henry (1926), "Principia Mathematica. Whitehead, Alfred North, Russell, Bertrand", *Isis*, vol. 8, no. 1, pp. 226-231, DOI: 10.1086/358383. (Cit. on p. 16.)
Sider, Theodore (2004), "Replies to Gallois, Hirsch and Markosian", *Philosophy and Phenomenological Research*, vol. 68, no. 3, pp. 674-687, DOI: 10.1111/j.1933-1592.2004.tb00375.x. (Cit. on p. 50.)
— (2011), *Writing the Book of the World*, Oxford University Press. (Cit. on p. 49.)
Siebel, Mark (2005), "A Puzzle About Concept Possession", *Grazer Philosophische Studien*, vol. 68, no. 1, pp. 1-22, DOI: 10.1163/18756735-068001001. (Cit. on p. 51.)
Smiley, Timothy (1995), "A Tale of Two Tortoises", *Mind*, vol. 104, no. 416, pp. 725-736, DOI: 10.1093/mind/104.416.725. (Cit. on p. 13.)
Snyder, Eric, Stewart Shapiro, and Richard Samuels (2018), "Cardinals, Ordinals, and the Prospects for a Fregean Foundation", *Royal Institute of Philosophy Supplement*, vol. 82, pp. 77-107, DOI: 10.1017/S1358246118000176. (Cit. on p. 59.)

Soysal, Zeynep (2020), "From Metasemantics to Analyticity", *Philosophy and Phenomenological Research*, vol. 103, no. 1, pp. 57-76, DOI: 10.1111/phpr.12695. (Cit. on p. 50.)

Stairs, Allen (2006), "Kriske, Tupman and Quantum Logic: The Quantum Logician's Conundrum", in, *Physical Theory and its Interpretation: Essays in Honor of Jeffrey Bub*, ed. by William Demopoulos and Itamar Pitowsky, Dordrecht: Springer Netherlands, pp. 253-272, ISBN: 978-1-4020-4876-0, DOI: 10.1007/1-4020-4876-9_12. (Cit. on p. 3.)

— (2016), "Could logic be empirical? The Putnam-Kripke debate", in, *Logic and Algebraic Structures in Quantum Computing*, ed. by Jennifer Chubb, Ali Eskandarian, and Valentina Harizanov, Lecture Notes in Logic, Cambridge University Press, pp. 23-41. (Cit. on p. 3.)

Steinberger, Florian (2015), "How Tolerant Can You Be? Carnap on Rationality", *Philosophy and Phenomenological Research*, vol. 92, no. 3, pp. 645-668, DOI: 10.1111/phpr.12165. (Cit. on p. 80.)

Susanszky, Peter (2023), "Trying to Adjunct Without Knowing How: Adjunction and the Adoption Problem", *Analysis*, vol. 83, no. 2, pp. 277-284, DOI: 10.1093/analys/anac055. (Cit. on p. 2.)

Szymanik, Jakub, Arnold Kochari, and Heming Strømholt Bremnes (2023), "Questions About Quantifiers: Symbolic and Nonsymbolic Quantity Processing by the Brain", *Cognitive Science*, vol. 47, no. 10, e13346, DOI: 10.1111/cogs.13346. (Cit. on p. 86.)

Tarski, Alfred (1943), "The Semantic Conception of Truth and the Foundations of Semantics", *Philosophy and Phenomenological Research*, vol. 4, no. 3, pp. 341-376, DOI: 10.2307/2102968. (Cit. on p. 46.)

Tennant, Neil (1987), *Anti-Realism and Logic: Truth as Eternal*, New York: Oxford University Press. (Cit. on pp. 70, 72.)

Thomasson, Amie L. (2014), *Ontology Made Easy*, New York: Oxford University Press. (Cit. on p. 50.)

Thomson, J. F. (2010), "What Achilles Should Have Said to the Tortoise", in *Thinking about Logic: Classic Essays*, ed. by Steven Cahn, Taylor & Francis. (Cit. on p. 14.)

Tomasello, Michael (1999), *The Cultural Origins of Human Cognition*, Harvard University Press. (Cit. on p. 74.)

Tomasello, Michael and Malinda Carpenter (2005), "Intention Reading and Imitative Learning", in *Perspectives on Imitation: From Neuroscience to Social Science - Volume 2: Imitation, Human Development, and Culture*, ed. by Susan Hurley and Nick Chater, The MIT Press, DOI: 10.7551/mitpress/5331.003.0007. (Cit. on pp. 83, 84.)

Valaris, Markos (2014), "Reasoning and Regress", *Mind*, vol. 123, no. 489, pp. 101-127, DOI: 10.1093/mind/fzu045. (Cit. on p. 23.)

vanMarle, Kristy (2018), "What Happens When a Child Learns to Count? The Development of the Number Concept", in *Naturalizing Logico-*

Mathematical Knowledge, ed. by Sorin Bangu, Routledge, pp. 131-147. (Cit. on p. 64.)

Von Fintel, Kai and Sabine Iatridou (2017), "A modest proposal for the meaning of imperatives", in *Modality Across Syntactic Categories*, ed. by Ana Arregui, María Luisa Rivero, and Andrés Salanova, Oxford University Press, pp. 288-319. (Cit. on p. 6.)

Warren, Jared (2015), "Talking with Tonkers", *Philosophers' Imprint*, vol. 15, no. 24, http://hdl.handle.net/2027/spo.3521354.0015.024. (Cit. on p. 55.)

— (2018), "Change of Logic, Change of Meaning", *Philosophy and Phenomenological Research*, vol. 96, no. 2, pp. 421-442, DOI: 10.1111/phpr.12312. (Cit. on p. 44.)

— (2020), *Shadows of Syntax: Revitalizing Logical and Mathematical Conventionalism*, Oxford University Press. (Cit. on pp. 3, 16.)

— (2021), "Defending Understanding-Assent Links", *Synthese*, vol. 199, no. 3-4, pp. 9219-9236, DOI: 10.1007/s11229-021-03200-w. (Cit. on pp. 36, 39, 42, 47.)

Weiskopf, Daniel A. (2008), "The Origins of Concepts", *Philosophical Studies*, vol. 140, no. 3, pp. 359-384, DOI: 10.1007/s11098-007-9150-8. (Cit. on p. 78.)

Weiss, Bernhard (2022), "From Tools to Rules: The Evolution of Rule Following", *Philosophical Topics*, vol. 50, no. 1, pp. 55-82, DOI: 10.5840/philtopics20225014. (Cit. on p. 85.)

Wierzbicka, Anna (1996), *Semantics: Primes and Universals*, New York: Oxford University Press. (Cit. on pp. 9, 69.)

Wikforss, Åsa (2009), "Are There Understanding-Assent Links?", *The Baltic International Yearbook of Cognition, Logic and Communication*, vol. 5, DOI: 10.4148/biyclc.v5i0.289. (Cit. on pp. 36, 38.)

Wilkinson, Eric (2025), "Logical Rationalism", *Journal of Philosophical Logic*, DOI: 10.1007/s10992-025-09796-5. (Cit. on p. 16.)

Williamson, Timothy (1996), "Knowing and Asserting", *Philosophical Review*, vol. 105, no. 4, pp. 489-523, DOI: 10.2307/2998423. (Cit. on p. 82.)

— (2000), *Knowledge and its Limits*, Oxford University Press. (Cit. on p. 37.)

— (2013), "Logic, Metalogic and Neutrality", *Erkenntnis*, vol. 79, no. 2, DOI: 10.1007/s10670-013-9474-z. (Cit. on p. 26.)

— (2017), "Semantic Paradoxes and Abductive Methodology", in *Reflections on the Liar*, ed. by Bradley P. Armour-Garb, Oxford University Press, pp. 325-46. (Cit. on p. 16.)

— (2022), *The Philosophy of Philosophy (2nd edition)*, Wiley-Blackwell. (Cit. on pp. 36, 41, 43, 44, 46, 48, 88.)

— (2024), "Accepting a Logic, Accepting a Theory", in *Saul Kripke on Modal Logic*, ed. by Yale Weiss and Romina Birman, Springer, pp. 409-433. (Cit. on pp. 27, 76.)

Woods, Jack (2019), "Logical Partisanhood", *Philosophical Studies*, vol. 176, no. 5, pp. 1203-1224, DOI: 10.1007/s11098-018-1054-2. (Cit. on p. 19.)
— (2023), "A Sketchy Logical Conventionalism", *Aristotelian Society Supplementary Volume*, vol. 97, no. 1, pp. 29-46, DOI: 10.1093/arisup/akad003. (Cit. on p. 16.)
Wright, Crispin (1986), "Inventing Logical Necessity", in *Language, Mind and Logic*, ed. by Jeremy Butterfield, Cambridge University Press, pp. 187-209. (Cit. on p. 15.)
— (2000), "Neo-Fregean Foundations for Real Analysis: Some Reflections on Frege's Constraint", *Notre Dame Journal of Formal Logic*, vol. 41, no. 4, pp. 317-334, DOI: 10.1305/ndjfl/1038336879. (Cit. on p. 59.)
— (2004a), "Intuition, Entitlement and the Epistemology of Logical Laws", *Dialectica*, vol. 58, no. 1, pp. 155-175, DOI: 10.1111/j.1746-8361.2004.tb00295.x. (Cit. on p. 16.)
— (2004b), "Warrant for Nothing (and Foundations for Free)?", *Aristotelian Society Supplementary Volume*, vol. 78, no. 1, pp. 167-212, DOI: 10.1111/j.0309-7013.2004.00121.x. (Cit. on p. 39.)
— (2007), "Rule-Following Without Reasons: Wittgenstein's Quietism and the Constitutive Question", *Ratio*, vol. 20, no. 4, pp. 481-502, DOI: 10.1111/j.1467-9329.2007.00379.x. (Cit. on p. 30.)
— (2018), "Logical Non-Cognitivism", *Philosophical Issues*, vol. 28, no. 1, pp. 425-450, DOI: 10.1111/phis.12132. (Cit. on p. 30.)