Towards a Cognitive Foundation of Mathematics

Thesis for the degree of

"Doctor of Philosophy"

by

Aviv Keren

Submitted to the Senate of the Hebrew University of Jerusalem

October 2018

This work was carried out under the supervision of:

Prof. Carl Posy

and

Prof. Aaron Sloman

Acknowledgements

I have moved to the Hebrew University not only to make my obsession with logic formal, but to embrace my ever-widening yet ever-specializing intellectual passions.[1] Little did I know how well the university would serve me and how far it would take me. The university's Center for the Study of Rationality in particular, has become my academic home, supporting me financially, nutritiously, intellectually, socially and spiritually. As for the topics at the very heart of this work, both mathematical logic (and mathematics at large) and cognitive science were actually thriving and passionate graduate-level environments in which I took part. Last, but no less evident, is the philosophical environment, with the official legitimacy it granted for me to pursue my long-time intellectual calling.

Pursuing my own intellectual calling, however, my way, required much more than that. It required someone who could take me under their wing, and lead me sure-handedly into this new and different field, straight to a PhD work that would combine my other backgrounds (which to me are integral) and suit my ambitions (while restraining or facilitating them as needed). Carl Posy took it upon himself to make a decent philosopher out of me, and the deficiencies still left on my part are in spite of his infinite meticulous efforts. He also took it upon himself to cultivate my own early vision together with me, into a mature research program, and to guide its pursuit to an impeccable standard (again, standing deficiencies notwithstanding). His contribution is fused into this work in more places and at more levels than I could list.

Later on, Aaron Sloman, whose work I've planned to get into anyway, independently found my research proposal online, and wrote to me that we have to talk. Two hundred lengthy emails later, he also joined in as a second advisor. Arriving late to the game, my work didn't always fit his views on what should be done and how. Yet, he painstakingly went through my many drafts, always striving to turn his own uncompromising standards into constructive suggestions. His influence on my more general agenda far transcends this work and the particular cases I was trying to make. I

[1] That this is not a contradiction is left as an exercise for the reader.

have found his scientific breadth, depth, independence, scope of vision and energies to be invaluable – and truly astounding.

At various stages I've discussed significant parts of my work in series of lengthy sessions with Stewart Shapiro. It was a real privilege to get to do that with someone whose many works, more than anyone's, I grew up on as a philosopher of mathematics. His input on my developing ideas was indeed as illuminating and useful as I could have hoped.

Hannes Leitgeb (who was a member of my PhD committee, alongside Oron Shagrir and Eli Dresner) invited me to LMU's marvelous Munich Center for Mathematical Philosophy (which has funded my trip). Being immersed for a month in such a high-quality academic group, with so many seminars, fellows and visitors who share my interests (and intellectual ideology) at this level of specificity, was invigorating, and provided an invaluable perspective on where my own work, which I've felt that I have failed to communicate sufficiently, thus requires further work.

Cognitive science is integral to this work and to my general program. Yet, working cognitive scientists are, by and large, too far apart from the philosophy of mathematics (and the background it necessitates) to enable engagement as deep as required. I was thus lucky enough to have Ran Hassin, Asael Sklar and Ariel Goldstien be willing to go into what I was trying to do (in the overly-challenging final chapter) and discuss it.

On a more personal level, I would like to thank, first and foremost, Einav Hart and Shir Ben-Or, especially for emotional support which has been crucial, in particular, for the progress and completion of this project, respectively. Ilan Nehama has also been invaluable throughout, as a rationality colleague and as a good friend (one among many in this intersection). As for my closest philosophy colleagues, Gil Sagi, Ran Lanzet, David Kashtan and Rea Golan, they have made the experience much more lively and friendly, while supporting and contributing to this work in various ways. Thanks go furthermore to my family, of course, and (in no particular order) to Amit, Avi, Itay, Jonen, Kfir, Klil, Lior, Maayan, Michal, Neta, Ofek, Ravit, Reut, Roee, Tal, Tom, Tau, W, Yoni, Ziv and others, for being part of my life. Last but not least, thanks to L, for no longer being part of my life.

Abstract

In this dissertation I approach the philosophy of mathematics and its foundations from the point of view of modern cognitive science. The first, introductory part, considers a cognitive foundations program, in general philosophical terms. The second part, which brings in the cognitive-mathematical substance, demonstrates how a cognitive approach can empower the exploration and provide insight on classical issues in the philosophy of mathematics.

The first part begins by introducing the general foundational context with respect to the cognitive revolution, motivating an approach which strives to ultimately bring together philosophy of mathematics with cognitive science. It then addresses the pressing anti-psychologistic concerns, which have traditionally been taken to render cognition irrelevant to the philosophy of mathematics. I dismantle these objections and argue for the possibility of a cognitive yet realist foundation, having all the properties that mathematicians standardly conceive of their subject matter as having. The first part closes with a review of some notable modern representatives of a cognitive approach to the foundations of mathematics: Stanislas Dehaene; George Lakoff & Rafael E. Núñez; and Giuseppe Longo.

In the second part of the dissertation (which can be read independently), I develop my own approach. This approach goes beyond the current, limited reach of mathematical cognition, to harness the scientific understanding of the general ways of the mind. I develop an abstract framework of object representation that makes room for mathematical objects on par with non-mathematical ones. This framework serves as a bridge, through which the vast scientific knowledge concerning the non-mathematical can guide our understanding of the mathematical. The general cognitive understanding is thus applied to the working mathematician's mind, suggesting a picture of what may be going on – to be later validated and/or updated scientifically. This strategy allows to reveal cognitive influences on the mathematics as perceived, and through this, shed light on mathematical ontology itself and bring us closer to it.

I demonstrate and pursue this approach in two case studies, each one revolving around a very particular point, central to the history of the philosophy of mathematics.

The first case study uses the framework to produce a novel account of the interaction between the ordinals and cardinals: How, in the finite, they come to merge into natural numbers and constitute them without our notice, whereas beyond the finite this interaction breaks down, forcing us to acknowledge their distinctness. This account exposes a pre-formalization, pre-mathematical influence of the cognitive machinery on the perception of the mathematical workings, which are artificially hidden by the standard formalizations.

The second case study takes the object-representation-based approach further to demonstrate how a naïve set theory course can naturally come to produce the illusory intuition behind the principle of comprehension. The first phase sketches a cognitive account of the student coming to master the course and represent its ontology. This sketch suffices for drawing some preliminary skepticism toward such an intuition (its discovered inconsistency aside). The second, constructive phase, puts into abstract terms a familiar cognitive phenomenon from visual perception, concerning the phenomenological unawareness of the underlying mediating cognitive processes with their riches and their limitations – the illusion of a detailed perception of the whole visual scene. The analogous phenomenon, as it would take place in naïve set theory, produces the distinct form of the illusory intuition.

Contents

Introduction

At the outer circles of the philosophy of mathematics, cognition can evidently be relevant, influencing mathematical practice. But it has not found its place at the foundation-centered core, which explores and strives to account for mathematics itself – its metaphysics, ontology, and (idealized) epistemology. Historically, in the philosophy of mathematics, psychology was rendered irrelevant by the anti-psychologistic arguments of Husserl and Frege, which have largely prevailed (as much as any arguments could, in philosophy). This has found, and keeps on finding, substantial support in the platonist working philosophy common to the practice – even more so with mathematics going into the infinite, growing away from mentally representable entities. A major problem for platonism and the central foundational movements (e.g. logicism) was that they tended to render the job that is left for the psychologists to pursue – of accounting for the connection between mathematicians and that mathematical reality they explore – mysterious and, largely, scientifically unapproachable.

The main alternative foundational interpretation of mathematics as mind-dependent, which came to be almost synonymous with this notion, was Brouwer's intuitionism. With Brouwer coming from within mathematics, the main technical substance of intuitionism was in reforming the whole of mathematics according to intuitionistically legitimate constructs and principles. But in the philosophical battle to win over these within mathematical practice, intuitionism has lost (leaving various technical contributions and making room for constructivity *within* now-standard mathematics). It was too restrictive and, well, too unintuitive by the vast majority of mathematicians (rejecting, in particular, logic's realist *law of excluded middle*). Any modern philosophy that strives to account for *modern* mathematics rather than dismiss it, would now have much more to take in for its subject matter (and reinterpret, if necessary). The question is whether this sociological failure of intuitionism was due to its particular brand of constructs and principles, or whether it testifies to the non-viability – in principle – of *any* mind-dependent foundational approach. This could depend on what "independence of the mind" is taken to mean. Brouwer's particular conceptualization aside, the very reliance on a *phenomenological*, a priori basis may have been the only reasonable, available cognitive option at the time. But from the modern scientific point of view, it is not only a particular and restrictive choice, but a problematic one.

Thus enters modern cognitive science. In general, mathematics aside, cognitive science explores the mental itself – not just behavior – in objective terms. It goes beyond how things appear to us, unveiling the mechanisms that control and bring about these appearances in their glorious computational richness and in their introspectively-surprising limitations and systematic biases. These are the mechanisms that orchestrate our interaction with objects at our base mental level.

The cognitive revolution has also overtaken and given rise to *mathematical cognition*, the subfield which explores humans' (and other animals') mental interaction with mathematical reality: from early (but systematic) developmental observations by Piaget; to behavior experiments that reveal, for example, natural abilities and surprising mental representations of numbers through approximation; to computational models that could account for such behavioral patterns; to the growing array of imaging techniques that hold promise for taking the science much further. The scientific field of mathematical cognition is a rich one – albeit rather limited in scope, as it is currently focused mostly on elementary mathematics, which is the most concrete part of mathematics and therefore the easiest to explore (and which potentially may be supported by its own dedicated, innate mechanisms). But the field's ultimate concern is with thinking – not with mathematics. Cognitive scientists do not ordinarily presume to tell mathematicians and philosophers of mathematics what "mathematics really is" (metaphysically) nor anything about its ontological structure.

As the now-traditional, widely accepted anti-psychologistic stance would have it, the mental interaction with mathematics is one-way; human thinking does not bear on mathematical reality (and, by association, on its professional exploration by mathematicians as well as philosophers). And indeed, although the cognitive revolution is by now old news, it has not left a mark on the philosophy of mathematics – and certainly not on its foundational explorations. (Moreover, joint professionalization, in cognitive science too, is demotivated). But from the perspective of cognitive science itself, it is anything *but* irrelevant. Trust in mathematicians' intuitions about what is self-evident, even if regimented by logical consistency, is cognitively naïve. Cognitive heuristics, approximations, and evolutionary pragmatics in general (which in particular cares only for statistics of fitness value rather than absolute truth) – are the general rule. In the face of these, underlying, non-conscious computational richness should be accounted for – particularly when one is after precision. More philosophically, the assumption of a non-physical reality is a non-starter if it does not provide cognitive science a handle on the exploration of human interaction with that reality. In fact, as a working philosophy, cognitive science is generally naturalistic, reducing all that is non-physical to the mental (which in turn is reduced to the physical-computational). This drives a common view among cognitive scientists, of mathematics as a creation of the mind or somehow grounded in it. This, as mentioned, is incompatible with the mainstream philosophy of mathematics. Accordingly, mostly in the past two decades, a few figures coming from a cognitively informed position have begun to attend explicitly to the *philosophy* of mathematics, considering the main lines of work there as cognitively naïve and suggesting revolutionary alternatives. Meanwhile, scholars from within the philosophy of mathematics may tend to consider these cognitive-based works as rather dismissive of the unique nature of mathematics and inattentive to the intricacies of the topic (see e.g. section 3.3.4).

I personally think both sides are correct in their critiques and find this situation unacceptable – and hopefully temporary. This situation calls for a program to delicately integrate cognitive science into the methodology, and cognition itself into the foundational picture. My overarching goal is, thus, to open a conceptual and methodological dialog between the two sides. Both may find that they need to change in order to adapt to a holistically-coherent scientific-philosophical picture. But for this dissertation – which is ultimately a work in the philosophy of mathematics – innovating on the side of cognitive science is not a goal but a necessary risk. I strive for it to be the least controversial in its background cognitive assumptions from this scientific field's perspective (as much as is possible for such a lively area of empirical, theoretical, and philosophical debate). Keeping to standard introductory sources where possible (such as textbooks and handbooks) provides some assurance against cherry-picking.

The dissertation, with its overarching goal, is divided into two parts. In the first part, I *motivate and consider a cognitive foundations program, in general philosophical terms*. However, I do not further pursue and develop the grand project of a metaphysically cognitive foundation directly in such terms. For the second part, I focus on the rich, cognitive-mathematical substance. I develop a methodological approach centered on the mental *representation* of mathematical ontology (as a proxy for mathematical reality) to make two rather particular philosophical points. The more modest goal of the second part, then, is to *demonstrate how a cognitive approach can empower the exploration and provide insight on classical issues in the philosophy of mathematics*. Hopefully, through further such work, a clearer foundational and metaphysical picture will emerge.

I hope that Part II appeals not just to philosophers of mathematics but working (if theoretically minded) cognitive scientists too. Thus, it makes for an author-approved partial reading, for which the text is generally adapted (barring various references to the general philosophical context).

The finer chapter structure is as follows:

In Chapter 1 – *Foundations*, I prepare the conceptual ground, motivate, and aspire to justify a cognitive approach to the foundations of mathematics. I do so through a conceptual-historical narrative. This narrative frames the logical foundations and their limitations in fitting terms for drawing parallels with behaviorism. It then considers the cognitive revolution (with respect to behaviorism and generally) and "applies" it to the logical foundations, to call for a cognitive turn in the foundations of mathematics too. This call concerns both a cognitive methodology and change of focus (like the one I pursue in Part II) as well as a cognitive metaphysics (Part I).

Chapter 2 – *Beyond Anti-Psychologism?,* concerns the very possibility of a cognitive approach to the foundations of mathematics, as called for in Chapter 1. I address, from the viewpoint of modern cognitive science, the old, familiar anti-psychologistic concerns that have been taken to render cognition *irrelevant*, and attempt to dismantle them. In particular, I argue for the possibility of a cognitive *realist* foundation, having all the properties that mathematicians standardly conceive of their subject matter as having (precise, objective, discovered rather than created, etc.) – except that it is not completely independent of the mind but, in a sense, grounded in it.

Chapter 3 – *Cognitive Approaches*, is a critical review of existing, cognition-based philosophical approaches to mathematics by a few notable figures. With Brouwer as the (phenomenological and very different) historical forerunner, these include such individuals as: Stanislas Dehaene as a central representative of the field of mathematical cognition (Dehaene, 1997); George Lakoff & Rafael E. Núñez, who have made the most systematic, modern attempt to ground (a considerable subset of) mathematics in cognition (Lakoff, et al., 2000); and Giuseppe Longo, whose writing on the topic is probably the most extensive (Bailly, et al., 2011) (Longo, 1999; 2002; 2005; 2005; 2007; 2010) (Longo, et al., 2010). Through the topics that come up and comparisons with my own view, I will expound further part I's theme of a cognitive foundations program.

Here ends the programmatic, introductory Part I.

Chapter 4 – *My Representational Approach*, opens Part II by introducing my own approach as I develop and pursue it for the rest of the dissertation. Briefly stated, it is *representation* based; it aims to reveal cognitive influences on mathematics as perceived, and through this, shed some light on mathematical ontology itself and bring us closer to it. It goes beyond the current, restricted scope of the field of mathematical cognition, to harness the scientific understanding of the *general* ways of the mind. The strategy is as follows, to:

1. Adopt some common existing principles, conceptualizations, and theories from cognitive science.
2. Reframe these in abstract-enough terms, to generalize over both the standard realms already attended to by the science and the mathematical topic of interest, in terms of its mental representation.
3. Adapt and apply these principles/conceptualizations/theories to the working mathematician's mind with regard to that topic, to suggest a picture of what may be going on. The picture should be such that a common (relevant) cognitive scientist could consider a reasonable first guess, later to be validated and/or updated scientifically.
4. Draw suggestive philosophical conclusions.

I discuss the philosophical status of such an approach and the methodological complexities and risks. Finally, I set up the general context in which my two case studies take place – set theory. For the dissertation, this focus has the advantage of an emergent critique of set theory's traditionally central foundational role. The added riches that may be in the interaction between other kinds of mathematical topics and cognition (e.g., connections between mathematical continuity and the vast cognitive support for the spatial) are left for other such expeditions to demonstrate and explore.

Chapter 5 – *The Architecture of Cognition*, brings in a selection of fundamental, familiar topics and conceptualizations from general cognitive science, relatively established and well explored, that seem to pertain to mathematicians' relevant cognitive workings as well: working memory, automaticity, perceptual hierarchy, and the place of language. These are in the background of Part II. Later developments of this work will need to take into account scientific updates and rival theories.

In Chapter 6 – *Cognitive Object-Constitution*, I develop a preliminary abstract framework of object representation (in terms of humans' *interaction* with objects) that makes room for the representation of mathematical objects on par with that of non-mathematical (mainly physical) ones. This framework serves as a bridge, through which the vast scientific knowledge concerning the latter can guide our understanding concerning the former. I describe the following:

1. The fundamentals regarding the system's treatment of types of objects (e.g. apples).
2. How such types can (in a technical sense) be "extended" (e.g. oranges as a kind of fruit that can be squeezed) and "combined" with others (e.g. bats as a kind of winged mammals).
3. How these issues depend on the environment.
4. What principles guide actual development – learning, mastering, and maintaining this sort of representation.

Some cognitive background material precedes the account as a context for what it aims to capture and abstract from; further cognitive material succeeds it to demonstrate how to evaluate the framework scientifically. While leaving, for future work, the many details required of a full framework, this suffices for supporting the point of the next case study.

In Chapter 7 – *Ordinals vs. Cardinals in ℕ and Beyond*, my first case study, I provide a novel account, in terms of my framework for cognitive object constitution, of the intricate interaction between ordinals, cardinals, and natural numbers, between the finite and the infinite. This account disputes the classic mathematical tale of what actual infinity supposedly did to our concept of numbers and substantiate a different view: From the very beginning, even in the finite, ordinals and cardinals are inherently – mathematically – distinct, different types of objects. The restricted domain of the finite, however, allows for merging the two and bringing about the compound objects that numbers in fact *are* and “forgetting” (consciously) about the finite ordinals and cardinals themselves and the interaction between them. But the finite is simply the implicit context in which these mathematical objects happen to first be experienced and discovered or taught. It is a contingent, empirical statistic rather than an absolute mathematical necessity – one that the cognitive system, nonetheless, by its nature, comes to process equivalently. The introduction of actual infinity, much later (in history and in studies), merely forces us to notice the notions' distinctness. Foundationally, this distinction is already at the heart of things – and of their mental representation too.

Chapter 8 – *The Universal Set in Naïve Set Theory*, describes a larger-scale case study, which takes the representation-based approach further in order to account for how an elementary course in naïve set theory could come to produce the illusory intuition behind the "principle of comprehension" (which infamously leads to Russell’s paradox). For its first phase, I sketch a cognitive account of a student coming to master the course and represent its ontology, and then draw some preliminary skepticism toward such an intuition (its discovered inconsistency aside). In the second phase, I suggest a *constructive* cognitive account, one that explains the particular form of the illusory intuition, of "the set of all sets" (and through it, comprehension). For this, I introduce in abstract terms a familiar cognitive phenomenon from visual perception concerning the phenomenological unawareness of the underlying mediating cognitive processes, with their riches and their limitations – the illusion of a detailed perception of the whole visual scene. Considering the analogous phenomenon, as it would take place in set theory, then produces the sought explanation.

Part I:

A Cognitive Metaphysical Foundation?

1 Foundations for Mathematics

(Linnebo, 2017) mathematical *platonism* to mean that:

I) Mathematical objects exist.

II) Mathematical objects are abstract.

III) Mathematical objects are independent of intelligent agents and their language, thought, and practices.

Philosophical sophistication aside, platonism is considered one of the mathematician's default working assumptions (Linnebo, 2017). But as a working assumption, it furthermore assumes (ideally) direct awareness of mathematical objects as they really are (unless mathematical developments come to force skepticism toward that direct "perception", a.k.a. "crisis"). *Naïve platonism* puts an ontological and metaphysical trust in how the objects appear to mathematicians. Consequently, it rejects any sort of metaphysical reduction of mathematical objects to, or grounding them in, or accounting in terms of, anything else. They are just as they appear to be, "sui generis", and stand on their own.

However, naïve platonism as a philosophical stance can be difficult to accept (as its prefix may suggest). Trusting such direct awareness immediately raises questions regarding what generalized sort of "senses" provide it, how they come to let us "see" into that non-physical and human-independent realm, and why we should therefore trust them. To simply accept naïve platonism and put all such problems aside as challenges for scientists and philosophers of the mind is unwarranted (though not uncommon).

Foundations begin where naïve platonism is not enough – where a theoretical framework for mathematics, underlying what is perceived, is required or could have something to contribute, at least. Such an underlying framework need not necessarily contradict *non-naïve* platonism (as with Frege's logicism). But when the mathematicians' intuiting of their own subject matter is no longer the final court, platonism loses this central support. The metaphysics of mathematics then becomes dependent on what that framework makes it to be, what is assumed to be in the background, deeper than the explicitly perceived.

Various frameworks have been considered for foundations:

> The most prominent, of course, is set theory. The received codification of this is the axiom system ZFC, but there are other set theories, as well as extensions of ZF with new axioms like $V = L$, determinacy, and large cardinal principles. Other proposed foundations, each with a corps of dedicated advocates, are higher-order logic, structuralism, traditional logicism, neo-logicist abstraction, proof theory, ramified type theory and category theory (Shapiro, 2004).

There are various senses of "foundations" and what they are for. (Shapiro, 2004) divides them roughly into three types:

> First, a ***metaphysical foundation*** reveals the underlying nature of mathematical objects… Secondly, an ***epistemic foundation*** reveals the true proofs or justifications for mathematical propositions, or provides one way in which mathematical propositions can become known…
> Thirdly, a ***mathematical foundation*** is a theory into which all mathematical theories, definitions and proofs can be translated, at least in principle. (My emphasis)

Focusing on the debate of set theory versus category theory (versus the multiverse), and perhaps somewhat more on the side of mathematical practice rather than the philosophical, (Maddy, 2017) distinguishes various foundational *goals*, as follows:

i. *Meta-mathematical Corral* – bringing together all of mathematics under one roof – for then "proving something general about classical mathematics".
ii. *Epistemic Source* – where "the problem of mathematical knowledge reduces to [e.g.] the problem of knowing the set-theoretic axioms".
iii. *Metaphysical Insight* – i.e., into the "nature" of mathematical objects and structures. (A consequence of *metaphysical foundation*, but conceivably not requiring such a full-blown account).
iv. *Elucidation* – through the "replacement of an imprecise notion with a precise one".
v. *Risk Assessment* – for various considered definitions (e.g. $\sqrt{-1}$), structures (e.g. the hyperreals), or axioms (e.g. relative consistency between large cardinals).
vi. *Shared Standard* – "of what counts as a proof" (i.e., sociologically).
vii. *Generous Arena* – "a single arena where all the various structures studied in all the various branches can coexist side-by-side, where their interrelations can be studied,

shared fundamentals isolated and exploited, effective methods exported and imported from one to another, and so on".

viii. *Essential Guidance* – "capture the fundamental character of mathematics as it's actually done… guide mathematicians toward the truly important concepts and structures".

I readily accept Maddy's assessment that set theory contributes significantly to *meta-mathematical corral*, *elucidation*, *risk assessment*, *shared standard*, and the *generous arena*, while not to *metaphysical insight*, *epistemic source*, or *essential guidance*. This indeed shows that "reduction doesn't, by itself, dictate any particular foundational use". Both she and Shapiro make the instrumentalist case that, given the different uses of foundations, there need not be only one foundation. In the face of set theory's admitted inadequacies, her defense of it amounts to "but it was not meant to do that":

For *metaphysical insight*, she says, "set-theoretic reductions don't give any sort of deep metaphysical information about the nature of the line or of ordered pairs or of natural numbers, nor are they so intended". Later, when considering a deep advantage of category theory, she takes this further to be an explicit argument: "I don't see that anything like *Essential Guidance* was among the ambitions of set-theoretic foundations in the first place, so to count them as 'failures' of set theory is to fault a cat for not being a dog".

Lack of intention may officially prevent an inadequacy from being a failure. But the intentions in the background are not what is philosophically important here, nor are they stable or well-defined enough to support such a role. The founder of some mathematical foundation might have had, to begin with, many hopes for what foundations in general would achieve. Some of them may have been abandoned from the start if they were obviously doomed to fail for that foundation. Setting $\{\{\emptyset\}\}$ as a surrogate for the number 2 is not (and probably never was) meant to suggest what that number actually, metaphysically is, simply because sets and numbers seem to be inherently different types of entities – with no further reason to suspect otherwise, at some deeper level. Intentions, then, were kept alive only for further uses that still could conceivably have succeeded. Frege's heroic attempt was metaphysically ambitious and failed in that, while ZFC had no such ambitions and so, for Maddy, does not fail. But this does not change the fundamental philosophical situation of the various possible foundations and their varied successes and failures within the space of foundational goals.

A liberal approach to the types of foundations and foundational goals might not seem objectionable. However, there is one foundation that I claim is, in an important sense, privileged: *metaphysical foundation*. Such a basis determines the makeup and structure of mathematical reality. Anything else is framed and defined, indeed *is*, in relation to it – even if implicitly and indirectly (given that there can be some foundational achievements even though we are not in possession of an ultimate metaphysical account yet). It determines the body of knowledge to which an *epistemic foundation* would hopefully systemize access. It constitutes the applicability of mathematics to other realms (an important foundational goal not on Maddy's list). And if a metaphysical foundation can itself be "mathematical" (like logic is), it is not only a *meta-mathematical corral* but the ultimate *generous arena*. Ultimately, the metaphysical foundation is the core of the full, unified foundation. I do not elaborate on these and the other foundations and goals. But consider again *essential guidance*, which was not "among the ambitions of set-theoretic foundations":

Maddy's suggestion "is that we do best to retain set theory in the foundational roles i[t] plays so well… but also pursue a serious philosophical/methodological investigation of the various 'ways of thinking' in mathematics". One could hardly object to pursuing a serious investigation of a serious matter. But keeping to the foundational (anti-psychologistic) division that is in the background here may ultimately prevent making progress on either front. "Ways of thinking" do not reflect simply how mathematicians psychologically explore the mathematical realm. They reflect the *interaction* between the mathematical realm and its exploration, which determines why some concepts, structures, proof methods, and the like are particularly important, for example. There may be some extra structure to mathematical reality itself that plays a role in that interaction (influencing the paths that mathematical practice takes within that reality), an extra structure that encoding into set theory may miss (and category theory perhaps does not, at least partially). Unlike the set-theoretic foundations, the true metaphysical foundation does also ground this extra structure and provides the arena in which to explore this interaction between mathematical reality and our exploration of it, psychological and otherwise.

The central status of the metaphysical foundation gives it prime importance and renders it as probably the hardest part to provide – the summit we conquer last (if at all). Along the way, there may come other foundations. There may be partial foundations (e.g. *epistemic foundation*) that capture one philosophical aspect of mathematics, even if we do not yet possess a complete picture to integrate them into (e.g. leaving the metaphysical foundation wanting). And there may be instrumental foundations (e.g. *meta-mathematical corral*, which need not be unique) that offer instrumental value even if they do not explicitly carry any greater philosophical significance or truth. But this does not entail an *ultimately* liberal foundational approach. There is a substantial interaction between the many issues explored. Not only would disparate foundations miss these interactions, but this may restrict them from achieving their own "intended" goals in full. In the end, I thus claim, we should strive – *holistically* – not for more foundations, but for foundational *unity*. And the metaphysical foundation is what would stand at the core of such a unified ultimate foundation of all and everything about mathematics (whether we can figure it out or not, perhaps not even in principle). The question is how to approach it, and less ambitiously, how to approach the aspects of mathematics for which current foundations are inadequate.

How can foundational research make progress? To begin, what can it build on?

Mathematics, going beyond its own platonic realm (if there be one) and materializing in ours, reflects itself and its structure in it. But naïve platonism, accepting a *direct*, *transparent* link between mathematical reality and mathematicians' perception of it, the link working flawlessly and fully, thus revealing mathematical reality's true, full nature, that nature having no other underlying, foundational structure – will not do. The only way to unveil any such structure, then, is to observe and explore mathematics' earthly manifestation, reflect upon the non-transparent link – and through it, upon mathematics.

Mathematics appears to us in various forms. It is expressed as follows:

i) <u>Implicitly</u>, in nature, for us to read mathematics from it. Not only in physical laws (etc.) but also in computer simulations or neurological mechanisms that do the information processing – and in their mental product, in which mathematics can play an *implicit* role. For example, it has a role in suspecting that, since only two of three bears left the cave, one might still be in it. (Not necessarily being sure about it, though, because not seeing mathematics *as* mathematics can entail not seeing its absoluteness either). Unless the bears were individuated in perception, when it is not known *which* bear might still be there, the reasoning is mathematics *based*. But without passing through numbers per se, abstractly, it is not *explicitly mathematical*.

ii) Explicitly but subjectively, when mathematicians consciously perceive and do mathematics in their heads *qua* mathematics. This can include errors – as these can still relate to mathematical reality in the broad sense (if imprecisely). For example, deeper mathematical subjects may be harder to grasp. This measure of depth may be mathematical par excellence (in a sense that we would like foundations to capture). But its pedagogic expression may be approached as such (whereas nature might simply not reflect levels of depth; all mathematics is reflected equally, absolutely). Let us rule out nothing a priori.

iii) Explicitly and objectively (observably), in the physical manifestations that interact with the subjective: calculations, equations, definitions, diagrams, informal scribblings, etc.; in speech, in (literal) hand waving, on the board, on a disk, in a rejected paper, or a published one.

The part of it done by professional mathematicians is sometimes systematically socialized (and roughly standardized) into so-called "mathematical practice". But from our point of view, *actual* mathematical practice includes not just the subjective (e.g. repeating an argument if the listener did not understand), but also the implicit (e.g. intuitions such as "a perception of beauty" that can lead mathematicians to focus on certain topics or be influenced by certain proofs, etc.) – implicit that may yet one day be made explicit, like proofs once were.

All these together are mathematics' manifestations (i.e., broadly construed); they are *all that is mathematical*. But there are certain requirements for using any of them as a foundational basis.

1) Somewhat trivially: The basis cannot itself simply *be* the whole riches of mathematics and its various manifestations. It needs to be more restricted or fundamental, so it could *give rise* to the rest in some sense. The opposing stand is *anti-foundationalism* (Putnam, 1967).
2) As theorizers, whatever manifestation we observe and try to build upon must already be explicit to begin with. Even if we observe and explore an animal's

implicitly mathematical reasoning, we already take it qua mathematics. Otherwise, we would not be able to designate it as (implicitly) mathematical. Mathematics' boundaries must be drawn first so that we can observe, explore, and reflect upon *that*. (Mathematics may certainly still expand beyond its current agreed confines, as it has done before. This may happen to show that the basis, too, was incomplete). This is the limitation of any theory we could relate to. To paraphrase Bishop: If God has a philosophy of mathematics of his own that needs to be done, let him do it himself.

3) The basis should be (at least) as precise as the mathematics which it is to ground. But earthly manifestations, or at least our grasp of them, may be – and, a priori, would be – imprecise. Historically, this was an important objection to grounding mathematics in psychology – a matter which Chapter 2 attends to at some depth.
4) The basis should not be tied *too* directly to any specific natural-world application or manifestation of mathematics – if mathematics itself is not (as, for example, was becoming clear in the 19th century with non-Euclidean geometry). Some chunk of mathematics developed with a specific application in mind, and allegedly grounded in whatever basis, might one day come to find a novel application; that basis should still ground it just the same.

Well into the 19th century, mathematics had yet to have any such basis on offer that could live up to these fundamental requirements. And then came modern logic (partly driven by the rapid development of mathematics itself), thereby opening the era of foundations, allowing for substantial progress beyond naïve platonism. Let us now briefly reflect upon it and its place in the history of foundations.

1.1 Logical Foundations

Within the variety of mathematics' manifestations, a central explicit, observable form in which it is expressed (especially in the professional's practice) is mathematical language. (So much so, in fact, that it is often claimed that "Mathematics is a language", simpliciter[2]). Definitions and mostly proofs, standing at the heart of mathematical language, may appear as almost constitutive of mathematical practice. Some

[2] To quote one prominent mathematician: "The basis of all human culture is language, and mathematics is a special kind of linguistic activity." (Manin, 2000 p. 154)

mathematics *can* and does also take place visually (sometimes even whole proofs) – but only occasionally and to a limited extent. Only mathematical language seems sufficiently general, intersubjective, and in fact objective, to be a *candidate* for reflecting the essentials of the practice and, through it, grounding all of mathematics. The hard-to-account-for and, frankly, mysterious connection between us and mathematical reality, then *factors* through language: Ontology is (somehow) connected with logical language, which in turn is (somehow) connected with us. Each might still be challenging on its own, but this option does seem to offer at least the potential for progress (for example, exploring the connection between us and propositions within this general, not-just-mathematical context).

Mathematical language does not simply come this way, as an organic product of mathematical practice that has grown out of the *use* of its practitioners. It had to be made explicit, as a well-defined system, rather than just an unsystematized background phenomenon (Ferreiros, 2001); only then could it be brought to serve mathematics as a foundation. Frege, Russell, Hilbert and others created modern mathematical language as a mathematical system of itself (grounded in formal logic), which could then serve as the framework where mathematical practice takes place (or be interpreted in, retroactively). Complementing this, Dedekind, Hilbert (again), Bourbaki, and many others integrated mathematical language as a framework into the very fabric of modern mathematics and its practice, reforming them around that abstract framework.

This can all be viewed as a general progression of mathematics (as a field) becoming more self-reflective and even self-conscious, taking less for granted, unexamined, the connections between the mathematicians and the reality they are exploring. Systematized reflection is what is at heart, whether it is done for purely mathematical motivations or for more philosophical ones. The interface not only provides a new realm of exploratory tools and methods, but it also reveals underlying order to the mathematical reality being explored. Philosophically, none of this *prohibits* holding on to platonism (which is indeed still held by many mathematicians and philosophers of mathematics). However, it does shift the focus from the ("directly") perceived ontology to the underlying means by which it is given to the mathematicians and through which they refer to it (which, as noted in the previous section, can weaken the appeal of platonism).

And so the *logical foundations* of mathematics came to be (a term that I use to designate both *logicism* (Demopoulos, et al., 2005) and *formalism* (Detlefsen, 2005) when what is important is the technical nature of the basis rather than the metaphysics attributed to it). They took on a central foundational role, to a considerable extent as the "official" foundations of mathematics, a position which they still generally hold. An important, modern reason for this is their undisputed *instrumental* value (as elaborated upon and granted in the previous section).

Mathematical motivations bred *mathematical logic,* which grew from logic into a rich mathematical field that can be devoid of (explicit or even implicit) philosophical content. Logic as a foundation for mathematics provides not only inherent mathematical *interest* for its own sake, but the ability to serve mathematics at large, mathematically (i.e., by illuminating mathematical concepts or even proving mathematical results through application).

The early adoption of the logical foundations, in the first decades of the 20th century, was also due to their *promise* of being *the* (full-blown) epistemic foundation of mathematics and, in some variants, even its metaphysical foundation. This promise has collapsed in interesting ways (particularly for our topic), to which we now turn. However, this collapse has not deeply shaken the logical foundations' status – mostly due to the lack of competition or complementary foundations. They are still "the closest thing we have", the best approximation of what a foundation for mathematics would be, capturing at least an (important) part of it. They are still the most *instrumental* for what we would like to get out of foundations. *Foundations of mathematics*, as a modern *field*, is somewhat distinct from mathematical logic. Ultimately, it is driven by philosophical motivations; it sheds philosophical light on mathematics, even in the face of the collapsed general promise. However, the extent to which logic is its fundamental tool too, tends to blur this distinction tremendously.

The first logical foundations endeavor, to have infamously collapsed, was Frege's logicism (Zalta, 2018) in its original form, with its reliance on "Basic Law V" (as a self-evident, if not necessarily obvious, axiom). This "law" was shown by Russell to be contradictory (in the footsteps of Cantor's diagonalization). Fundamental principles that we can see to be self-evident were to be the safe basis on which to build everything else. That we cannot trust even *those* not to be outright contradictory (let alone true) is a crucial lesson for us here, in a sense undermining the very possibility of anti-psychologism (see Chapter 2).

Frege's ambitious brand of a logical foundation was also rendered unachievable at large, for mathematics as it has been extended by Cantor and others into the infinite. For Frege, not only $\mathbb{N}$, as a structure, was grounded in logic. Every natural number, as a mathematical object, was to be grounded in a logical object, individually. (Let us call a per-individual-object grounding in general ***object-complete***). But accepting an uncountable universe of objects no longer allows for this, under reasonable restrictions on the language. Lifting the restrictions and strengthening the logic (e.g. to an uncountable language or ontology of concepts to begin with) could technically solve that. However, it would be philosophically unmotivated, defeating the point of a worldly (if idealized) manifestation that humans can relate to, explicitly. Ontology has grown apart from epistemology (which is what logic is fundamentally about). Thus, from Russell and Hilbert onwards, logical foundations can no longer presume to be object-complete. They may *ground* theories and structures, but only certain individuated objects. They must therefore be *meta*-mathematical, leaving an inherent *gap* between the foundation's objects – and the mathematical ones.

Whitehead and Russell's own foundational endeavor that followed Frege's, the Principia Mathematica, gained much more immediate fame and influence, sparking the era of foundations for mathematics (alongside Zermelo's first axiomatization of set theory). Importantly for us, the reasons for its (correspondingly famous) failure, as well as Hilbert's (at which we arrive a little later), are still as relevant as ever. This is the starting point for consideration of *any* logical foundations for mathematics that may be. In general terms, it is about the limitations of logic (which depend on the specific type of logic).

Logic-based definitions serve as an important part of mathematicians' (linguistic) referential gateway to the mathematical realm – to its ontology and its structure of truth (as revealed through proof). As discussed, logical foundations privilege definitions and proofs, within all that is mathematics related, as what *grounds* mathematics. For that to succeed, for them to be *foundationally adequate* as our gateway to the realm, they must (as two central requirements) determine the ontology and the truth about it. As it turned out, they fail on both accounts – under the more standard understanding of logic and its constraints, which I address first.

That understanding, which emerged out of the (underspecified) early foundational explorations, settled on 1st-order logic as *the* logic for mathematics (and beyond). Its main advantage over higher-level logics (which are more expressive) is that its logical consequence relation is captured by finite provability (having (sound and) *complete* deductive systems)[3]. Thus, abiding by 1st-order logic reflects being serious about proofs, too (rather than just definitions and semantics), as *foundationally constitutive* of the structure of truth. The alternative (i.e. higher-order or infinitary logic) sets proofs aside (I think is fair to say) as more of an epistemological tool for mathematicians' *exploration* of the truth of that reality, which need not be part of its foundation. Within the general move of privileging of all that is mathematical for foundations, this alternative would deplete its linguistic part of much of its structured substance.

Accepting the role of (finitary) proofs as fundamental, however, comes at a price. Completeness entails (the intimately related) compactness and (for various theories and models) the existence of so-called "non-standard" models[4]. Focusing just on ℕ (as I do here), it entails that there are models with "infinitary numbers", not isomorphic to ℕ – that uphold the very same (*first*-order) laws that hold with respect to the "real" natural numbers.

For a foundational approach that enshrines proofs, this does not *need* to be a problem; proofs pull through just the same. If they also thereby happen to track the truth about models other than the intended one, then so be it. After all, this is generally a feature of the absoluteness of proofs rather than a bug, pertaining beyond whatever whoever set them up had in mind (allowing, for example, the development of group theory with some disregard for *any* particular model of it). In reality, mathematical *practice* can never run into this phenomenon: As long as it all takes place within some well-defined language, such elementary-equivalent yet non-isomorphic differences are "metaphysically-speculative extra luggage". They could never come to present

[3] For further details, see any mathematical introduction to 1st-order logic and its metalogical properties (and Lindström's characterization of it in terms of those properties).

[4] See (Gaifman, 2004) for an overview.

themselves, causing a disagreement about whether some *exhibited* number is really a natural number or not.

Yet this assurance seems insufficient. Mathematics, by design, abstracts structures from differences that are up to isomorphism (Shapiro, 1997) and takes as meaningful all differences that are left standing. Logical-linguistic expressibility of differences seems too weak, overly "anthropomorphic", subjective, not pertaining to the mathematical reality as it is – independently of some linguistic expression of it. Non-standard models, then, open the door to the important challenge of successful *reference* to the intended model, paradigmatically to ℕ. The logical foundations were supposed or would have been hoped to account for this (not just for Frege but through Peano's axioms and the like). In actuality, mathematicians certainly do not come to be acquainted (at childhood) with ℕ through any such axiomatic fashion. Quite the opposite, they come to value and focus on certain axioms or laws *because* these seem to be true particularly of ℕ. ℕ itself appears to be given to them (to us) in a wholly different manner, somehow. But it does seem that they all have that same distinctive structure in mind, well, "well-defined", a pure mathematical percept par excellence – if anything is. If logical foundations cannot, even in principle, offer a way to designate or specify ℕ, if they offer no such control over the connection between the linguistic level and the semantics content we need for it to be about (which does not seem like a "merely" psychologistic concern), then it cannot be the case that mathematicians access mathematical reality exclusively through them; they cannot be the interface that fully grounds everything we take to be truly mathematical.

To be sure, the focus on ℕ here is quite essential. A realistic belief in *the* (i.e. unique) universe of sets would mean that mathematicians would like to be referring to *it* when they do axiomatized set theory, rather than some non-intended, non-standard model. However, unlike for ℕ, it is an unreasonable expectation to always have in mind the very same *set* model. Unlike ℕ, these are creatures too infinitely complex (and high) to entirely see with mathematical clarity in the mind's eye in *any* sense. Mathematicians have only a very partial representation (at best) that leaves much underdetermined. One may justifiably feel uncomfortable in the face of a *countable* model of ZF (containing, as always, "uncountable" sets). But in general, realism toward unique reference to the set universe seems to be of a much different degree (if not a wholly different character) than just unique reference to ℕ (for which biting the solipsistic bullet is much harder and almost never done). In fact, such a split attitude can be found there from the very beginning:

> "Non-standard models were introduced by Skolem, first for set theory, then for Peano arithmetic. In the former, Skolem found support for an anti-realist view of absolutely uncountable sets. But in the latter he saw evidence for the impossibility of capturing the intended interpretation by purely deductive methods." (Gaifman, 2004)

Beyond ℕ and set theory, it is rather customary to take the "intended" issue pragmatically, qua non-essential intentions, which raises no problem at all. Rather, compactness is a useful tool for producing new (but elementary-equivalent) structures that are of independent interest and value (e.g., hyperreals for non-standard analysis, pseudo-finite fields). But still: ℕ alone is central enough (for mathematics itself and in the history of its foundations) that the conundrum of the reference to it is an important, if not dire, problem (which is in the background of Chapter 7). When one insists on grounding the natural numbers in mathematical language alone for its foundation, an essential component is left missing.

So much for non-standard models as a necessary evil of 1^{st}-order, proof-grounded logic.

Even when the logic is properly finitistic in this proof sense, there is a further fundamental, far-reaching, infinitary idealization involved with the mathematical language itself, if it is to ground a structure such as ℕ. The perception, realistically conceived, is of a structure as *having all its truths to it* – truth-value realism (Linnebo, 2017). It may be argued that the practice reflects this in its optimistic directedness toward conjectures and open problems (as opposed to just growing the body of knowledge bottom-up, whatever it may find). But the set of truths is not given to begin with. It is reached, gradually and necessarily always partially, through deductions from a much simpler, *essentially finite* (recursive, effective) axiomatic base. This base (and particular deductions from it) are supposedly what mathematicians relate to. This is their interface, what foundations would need to ground mathematics in (rather than just grounding mathematics in something else, more idealized, that we cannot relate to). As Gödel famously showed, it turns out that such a base cannot suffice even for grounding the (whole) truth about a structure such as ℕ (which in turn does not fully determine the structure itself, as discussed above). One of the most basic structures in mathematics – which was there from its very beginning, and that can successfully be mastered by practically all humans – is too complex, in this important sense. Importantly, other structures, including seemingly more complex ones (such as the first-order theory of the real numbers), need not succumb to this. But with ℕ being so pervasively embedded throughout mathematics (rendering the theories that embed it incomplete as well), the phenomenon is quite general (especially with the modern understanding of it, following Turing, in terms of the limitations of computation).

Within the same thrust, Gödel also demolished Hilbert's meta-mathematical hope[5] of securing all of mathematics upon a finitary logical basis (Zach, 2016): A theory such as PA (if consistent) cannot prove even its own consistency – let alone that of a stronger system such as ZF. To do so (speaking metaphorically) would be like paying a loan by getting the money through taking another loan of a higher interest rate. We cannot receive more than we put in.

These foundational failures have been discussed extensively in the literature – which I hardly touch upon. They have given rise to updated (and usually less-ambitious) foundational approaches – which I do not touch upon at all – except for the following:

Still within the confines of *logical* foundations, the situation may seem less dire when one is willing to abandon the centrality of having (finitary) deductive proof systems and hence of first-order logic. (Shapiro, 2000) makes the case that 2^{nd}-order logic, unabashedly, is reasonable as a foundation of mathematics from a "model-theoretic" rather than a "proof-theoretic" point of view. With a similar attitude, (Hintikka, 1996) attempts to replace it with his (and Sandu's) "Independence-Friendly 1^{st}-order logic". This logic, though much more expressive than classical 1^{st}-order logic, has quantifiers ranging over elements only, not sets of elements, and shares some central properties with the classic. However, as (Feferman, 2006) points out, in another important sense, this alternative is still entirely 2^{nd}-order logic (a "blame" which Hintikka was trying to avoid). All this raises (fascinating) issues in the philosophy of logic, concerning why

[5] As originally conceived; See (Zach, 2007).

1^{st}-order logic gained its central status for mathematics in the first place (Ferreiros, 2001). A few points to take from this debate (which I do not enter into) are:

1) The matter of grounding mathematics in its language, its logic, cannot really be separated from that of the choice of logic itself.
2) There can be valuable substance to a foundation even if it fails to account, epistemologically, for mathematical *proof*.

To put this line of foundational approach into our context: When standard, "proof-theoretic" logical foundations strive to ground mathematics (rather than secure it, like for Hilbert), they try (or rather tried) to ground it, within all that is mathematical, onto mathematical language and its (deductive) structure of truth alone. This is a very slim basis to begin with. Shapiro's and Hintikka's "Model-theoretic" approach then takes the deductive structure, too, out of the equation. This leaves only language itself, as a means of expression that mathematicians use to refer to laws and structures. Going higher-order, they can indeed specify ℕ uniquely through its Peano axiomatization (non-intended models ruled out by induction now ranging over what it was intended to, not restricted to first-order formulas). They can express – directly through the logic – many other concepts needed in mathematics, concepts that standardly require set theory for background. Language here serves – successfully – as a *means of communication* between the mathematicians. They understand each other through it and can go on with their business (fitting Shapiro's and Hintikka's "mathematicians first" approach, see p. 38). However, our foundational business is a different one. How could that language alone – in principle – come to designate whatever *the mathematicians see* it to be about? Remember, anything occurring in their heads is not itself part of this slim foundational basis. The set of well-formed formulas is just that – and all that there is. Yet it somehow needs to be related, metaphysically, to the ontology of mathematics (which in itself it certainly is not). What deductions have given us is a "Functor", *tracking* mathematical truth: the *relation* between real-world sentences (in the mathematician's mind, on the board, etc.) transferred (metaphysically) to a *relation* between propositions in the mathematical realm. This is how that part of actual practice reflected something (quite a lot, actually) about the structure of mathematical reality (enough to be taken to ground mathematics). With that gone, foundations upon non-deductively-complete logics implicitly count on more than is explicitly there; they count on our ability to bestow meaning upon the language's expressions. Shapiro and Hintikka may not mind counting on this, for their own foundational purposes – but for mine, I do.

Making this rather deep point properly is beyond the scope of this introductory chapter. What is important to take from this discussion is that:

3) Successful communication alone is not enough for a full specification, a full grounding, as it builds upon implicit substance (possibly a lot of it) that is shared between the *communicators*.

This concludes the overview of the logical foundations as the central foundational movement, from the point of view I am developing. The critique, in itself, is similar to the intuitionist critique of the logical foundations, and certainly in its spirit. Where my narrative now diverges is in the *justification* for this point of view (and the alternative this suggests). For this, we first need to understand the place of the cognitive revolution in psychology – beginning with what it replaced.

1.2 Behaviorism

In psychology, in the first half of the 20th century, the doctrine of *behaviorism* gained central ground (mostly in North America) and rose to fame. It is the main contention of this section that *the logical foundations can be seen as a form of behaviorism with respect to mathematics*. I now present its conceptual basics (following the introductory (Graham, 2017)) while drawing parallels with the logical foundations as presented above. (There are also disanalogies, which are beside the point of the narrative).

Behaviorism involves "empirical constraints on psychological state attribution… insist[ing] on confirming 'hypotheses about psychological events in terms of behavioral criteria'" (Graham, 2017). The full doctrine (which includes, in particular, Skinner's famous *radical behaviorism*) is committed to the following three claims:

1B) "Psychology is the science of behavior. Psychology is not the science of mind -- as something other or different from behavior" (ibid.).

To start spelling out the rough analogy:

1L) Foundations are to ground mathematics through (an idealization of) mathematicians' mathematical behavior. They do not need to ground what is in their mind as something other or different from that behavior.

Logical foundations then build specifically on a central component of mathematicians' *linguistic* behavior, which, importantly: A) Lends itself to mathematical systemization (unlike most other facets of mathematicians' behavior). B) Is rich enough to potentially give rise (in some sense), through axioms and deductions, to all the rest of their mathematical behavior. That is, the assumption is that e.g. diagrammatic behavior is related to the logical-linguistic behavior in such a way that it can be interpreted as being *about* (e.g. an aid for) the latter, in a sense that captures what is mathematically essential.

2B) "Behavior can be described and explained without making ultimate reference to mental events or to internal psychological processes. The sources of behavior are external (in the environment), not internal (in the mind, in the head)". (ibid.)

2L) The foundational basis is to be found in the explicit that is *objective*, *observable*, not in the *subjective* (p. 14); ultimately, there is no reference to mental events or to internal psychological processes. The sources of mathematicians' (idealized, error-free) logical-linguistic behavior are external (in the environment), not internal (in the mind, in the head). Here, the environment is the axioms and deduction rules, and by mediation, the platonic reality with which the mathematicians are dealing, if one is postulated.

3B) "In the course of theory development in psychology, if, somehow, mental terms or concepts are deployed in describing or explaining behavior, then either (a) these terms or concepts should be eliminated and replaced by behavioral terms or (b) they can and should be translated or paraphrased into behavioral concepts" (ibid.).

3L) In the foundational theory, if mental terms or concepts are somehow deployed in describing or explaining mathematicians' behavior, then either (a) these terms or concepts should be eliminated and replaced by logical-linguistic terms or (b) they can and should be translated or paraphrased into logical-linguistic concepts.

An example for (a) could be eliminating the notion of *intuition* (that some mathematical proposition is true) from a foundational account and its replacement with (the requirement for) *proof*. This replacement rules out the legitimacy of mathematical knowledge that the brain, in its own ways, may acquire but that is not explicitly constituted by a formal proof; e.g., the inseparability of two connected rings (an example of Sloman). Here, logical foundations merely make formal an attitude that is central to mathematical practice to begin with.

The commitment to (b) is itself the partial doctrine of "'Analytical' behaviorism (also known as 'philosophical' or 'logical' behaviorism)". This is "a theory within philosophy about the meaning or semantics of mental terms or concepts. It says that the very idea of a mental state or condition is the idea of a behavioral disposition or family of behavioral tendencies, evident in how a person behaves in one situation rather than another" (ibid). On our reading, the very idea of mathematical objects being the idea of laws about them that define how they behave – is the hallmark of *formalism*. A central appeal of analytical behaviorism is that it "helps to avoid substance dualism… the doctrine that mental states take place in a special, non-physical mental substance (the immaterial mind)" (ibid.). Similarly, formalism is the standard way to avoid commitment to a separate, non-physical mathematical ontology (which the mathematician seemingly intuits) and all the metaphysical hurdles that come with it.

It is important to stress, however, that what is at stake is not "merely" "a theory *within philosophy* about the meaning or semantics of mental terms or concepts". Instead, "behaviorism is a doctrine -- a way of doing psychological science itself" (ibid.). To put the sub-statement (3-b) more elaborately: "Most … behaviorists and neobehaviorists… allowed for the introduction of theoretical terms, or postulated 'intervening [i.e. 'mental'] variables,' so long as these were, at least ideally, rigidly and exhaustively defined operationally, via principles or 'laws' relating stimulus inputs to internal states and internal states to behavioral outputs" (Greenwood, 1999). Thus, metaphysics aside, it is a restriction on the *form* of the theory. And the logical form, formalism shares with logicism. The mental reference to ℕ, which transcends (finite) observable behavior, need not be eliminated. However, at least ideally, it is *rigidly and exhaustively* defined operationally, via logical principles or laws relating it to mathematicians' input and output, of mathematical language. Logical foundations give this linguistic expression a constitutive status, regardless of what produces or drives that form of expression. Restricting mathematics' foundational basis this way is thus, on my suggested reading here, a "form of behaviorism" – and their failures may be related as well and are worth examining on par:

"The deepest and most complex reason for behaviorism's decline in influence is its commitment to the thesis that behavior can be explained without reference to non-behavioral mental (cognitive, representational, or interpretative) activity" (ibid.). "We don't just run and mate and walk and eat. We think, classify, analyze, and theorize. In addition to our outer behavior, we have highly complex inner lives, wherein we are active, often imaginatively, in our heads, all the while often remaining as stuck as posts, as still as stones" (ibid.). If this is reasonable for psychology at large, then it is even more so for mathematics as explicitly expressed in the world, not only through us but *in* us. Mental exploration and learning at large cannot be reduced to any simple environmentally defined form. Historically, there were good reasons for the domination of both behaviorism and the logical foundations – partly shared reasons. However, the tides have turned for psychology, and the reasons for *that* – may also be relevant to the logical foundations that still reign. Let us now turn to those changed tides in psychology.

1.3 The Cognitive Revolution

Behaviorism was dethroned (rather abruptly) by what has come to be known as *the cognitive revolution*. That "revolution" was primarily with respect to behaviorism, a paradigm shift that brought back several themes that had been already abundant in psychology's history and outside of North America. This is the conceptual side that the previous section spoke of. Its importance to this project is mainly that, through the suggested analogy with logical foundations, it can illuminate a supposed "revolution" or progress, from that still-reigning foundational approach to… something else, in analogy.

But with the cognitive revolution, something new also started, building especially on the new field of computer science and engineering, along with parallel advances in neuroscience. The cognitive revolution formed from a synthesis of research in various fields, with far greater interdisciplinary communication. To quote a leading figure in the revolution:

> "I argued that at least six disciplines were involved: psychology, linguistics, neuroscience, computer science, anthropology and philosophy… These fields represented, and still represent, an institutionally convenient but intellectually awkward division. Each, by historical accident, had inherited a particular way of looking at cognition and each had progressed far enough to recognize that the solution to some of its problems depended crucially on the solution of problems traditionally allocated to other disciplines" (Miller).

This attitude has been and still is growing (despite general pressures for professionalization) into *modern cognitive science* (or *sciences*, when one wishes to emphasize differences within it). It is a body of many empirical findings in various old and new subfields, shifting theoretical conceptualizations, and some shared fundamentals. Despite the rapid developments, one thing seems clear: there is no rolling back to behaviorism, conceptually (though *neobehaviorism* still finds its more limited place *within* cognitive science). This is what has come of the cognitive revolution.

The cognitive revolution and its further developments have been just as relevant to *mathematical cognition* (as introduced in the Introduction and to be elaborated upon on p. 71). But anti-psychologism has kept the revolution's influence from the philosophy and its foundations. Certainly, this dogma was not simply kept for dogma's sake. Logic (also taken to be independent of humans) makes for an extremely clear and precise basis (p. 14), which does (with set theory upon it) account for much that is mathematical, achieving various goals to a considerable extent. Modern cognitive science, on the other hand, is currently in no such position to serve as an extremely clear and precise, spelled-out basis. The anti-psychologistic concerns are for later, but there is nothing puzzling about the philosophy of mathematics having missed the cognitive synthesis.

However, behaviorism (as a methodology and a foundational philosophy) had its reasonable motivations and justifications too. General scientific progress cannot be ignored. New knowledge can force us to re-evaluate its relevance to the foundations of mathematics, as embedded within the larger, holistic philosophical-scientific picture. It is the central contention of this project that the cognitive revolution itself bears deeply on the philosophy and foundation of mathematics. I will now give a general review of central motifs of the revolution, all of them relevant to the philosophy of mathematics.

1.3.1 Mechanistic Explanations

From the works of Gödel, Church, Turing, and others in the 1930s arose a unified, extremely general notion of *computation*. Over the following decades, it has matured into concrete, working machines that have replaced what were once *human* computers – and took it much, much further. Computer science (and engineering) conquered and continues conquering new grounds and vastly expanding the horizons of what (in principle and in practice) could be done by such mechanical means (while delimiting what cannot).

This endeavor clarified the sense of what it actually takes to conquer a computational task (and how difficult that may be). It also clarified the extent of what the brain does as a great many computational tasks. General psychological principles would no longer do. Top-level observations, such as the Gestalt laws regarding how we group our visual experience, are not enough. To describe an actual mechanism that gives rise to a phenomenon is very different from listing laws concerning its behavior. Being able to program (at least in principle) a fully-functional visual perception system from the bottom is a necessary stage on the way to understanding how vision truly works. Conquering *partial* computational goals within visual perception clarifies the depth of the challenges that remain (a process that is still ongoing). This is completely general; an inability to program some human (or others') capabilities is standardly not taken to indicate that they are de facto non-computable (let alone in principle). Rather, it is taken on as a standing challenge, which could and should be met. This is Artificial Intelligence, the engineering pillar of cognitive science and the revolution.

Computers being "the new craze" did not mean that the brain itself necessarily is one. There is a long history of analogous assumptions that have turned out incorrect. But there are deeper theoretical reasons to take computation as a candidate for being *the mechanism* of full-blown information processing (Greenwood, 1999; Piccinini, 2012; Chalmers, 1994). And meanwhile, on the scientific side, computation was entering at all levels.

At the bottom level, unveiling the brain's own mechanism as a biological one: Early scientific explorations, including the fundamental idea that neurons are the basic structural and functional units of the nervous system, have been around since the late 19th century. But the actual behavior of neurons as atoms of computation was only (beginning to be) figured out much later, starting that aspect of the cognitive revolution. In 1943, McCulloch and Pitts showed how a network of artificial (simplified) neurons could form a computational model, beginning the classical computational theory of mind (Rescorla, 2017). And then later, in 1952, came the Hodgkin-Huxley model:

> The Hodgkin-Huxley (HH) model of the action potential is perhaps the single most important theoretical achievement in modern neurobiology. It consists of a set of differential equations that describe neuronal "firing". The model, and the experimental work that led up to it… [established] a new framework for thinking about the electrical activity of neurons. (Levy, 2014)

Neuroscience thus mandates conceiving the mind and exploring it in *ultimately* computational terms.

At the other, top-end level of psychology, Chomsky revolutionized linguistics through computation. This went far beyond accounting for (some early aspects of) linguistic abilities in computational terms. He reconceived language itself as a computational model (of generative grammars). A computational structure was written on the very face of a psychological phenomenon – once it could be searched for as such.

In between these extremal levels, it has become standard practice to account for psychological phenomena in semi-computational terms. Even if its grounding in the actual neural computation is left lacking. And even if the high-level phenomenon itself (e.g., behaviorally or consciously) does not in particular suggest, a priori, some computational structure.

Overall, this *mechanistic, computational standard of explanation* has become the standard in cognitive science (though there are many debates about various standards in cognitive science, including non-mechanistic ones, and many findings and theories that are not about causal mechanisms per se (Chirimuuta, 2017)). The standard's position proved a successful one: It was more restrictive than various other (mostly implicit) standards in the history of psychology, giving it considerable explanatory value. And it was far richer than behaviorism's overly-restrictive operational definitions, bringing the richness of its subject matter within its reach. Mostly, in between these, it was transpiring to be the way the mechanisms really are; promising ultimately to bring together high-level cognitive psychology with low-level neuroscience. The cognitive revolution, rather than just a body of empirical findings and isolated theories, thus brought with it a theoretical revolution regarding what it means to account for a phenomenon in these areas.

1.3.2 Mental Ontology

In comparison to behaviorism, the enriched mechanistic standard was complemented by a liberated mental ontology that no longer had to be operationally defined, conceived through its behavioral manifestation[6]:

> It is suggested that the move from behaviorism to cognitivism is best represented in terms of the replacement of (operationally defined) “intervening variables” by genuine “hypothetical constructs” possessing cognitive “surplus meaning,”…
>
> Of course, it is obviously also desirable that… operational measures be employed in the empirical testing of scientific theories, but these two desiderata — determinate meaning and operational measures — are quite independent (and are generally recognized to be so in the natural and biological sciences).
> Operational measures of theoretical constructs such as “electron” or “short-term memory” do not in general determine the meaning of such theoretical constructs. (Greenwood, 1999)

This shift came with the re-legitimatization of "new" types of theoretical entities, *representations*:

> a basic concept of the Computational Theory of Mind, according to which cognitive states and processes are constituted by the occurrence, transformation and storage (in the mind/brain) of information-bearing structures (representations)…
> [A] mental representation may be more broadly construed as a mental object with semantic properties. (Pitt, 2017)

(It is important to note that, although representations are ubiquitous and seemingly fundamental to the science, they are also a central topic of philosophical debate. For this preliminary project, I assume the mainstream *representationalism*). These entities and the computational mechanisms that process them (and thus the cognitive

---

[6] See (Sloman, 2008 p. 2024) for a different view on the limitation of behaviorism.

processes), the whole cognitive apparatus, became of central interest for its own sake. It was a shift that changed the very subject matter:

> [Cognitive scientists] have not merely aimed to develop cognitive explanations of behavior to replace traditional behaviorist… explanations of behavior. That is, from the point of view of cognitive psychologists since the 1950s, cognitive states and processes are as legitimate *objects* of theoretical explanation as are observable behaviors…
>
> [they] were as much (if not more) concerned with the explanation of cognitive processing as they were with the explanation of observable behavior. (Greenwood, 1999)

This brought in a whole set of concepts:

> While theoretical definitions of the sensory register, attention, long- and short-term memory, depth grammar, cognitive heuristics, visual perception, propositional and imagery coding, episodic and semantic memory, template-matching, procedural networks, inference, induction, and the like have abounded in the cognitive psychological literature, operational definitions — as opposed to specified operational measures — of these phenomena have been virtually non-existent. So there is a significant discontinuity between behaviorist and cognitive theories. (Greenwood, 1999)

For example, where it concerns Chomsky's approach:

> The grammatical rules that govern phrases and sentences are not behavior. They are mentalistic hypotheses about the cognitive processes responsible for the verbal behaviors we observe. (Miller)

All together (though this is not universally acceptable):

> a cognitive theory may be reasonably defined (following Fodor, 1991) as any theory that postulates representational states that are semantically evaluable—that can be characterized as true or false, or accurate or inaccurate—and rules, heuristics, or schemata governing the operation of such representational states, as they are held to be involved in receiving, processing, and storing information. (Greenwood, 1999)

1.3.3 The Non/Conscious

It is easy to take for granted cognitive abilities that everyone around us (including many animals) possess. The complexity that may underlie them, that there may be to the tasks themselves (that we handle naturally), is mostly revealed through the attempt and difficulty in *programming* those abilities (Section 1.3.1). The extent of this richness, combined with the related empirical exploration, brought about (as a continued discovery) a modern understanding, one that is quite different from our intuitive one:

Almost all that is going on under our hood takes place not only subconsciously but completely out of introspection's potential reach. The tasks that are the most important and frequent are handled automatically, hidden from consciousness so to not consume its conscious resources. These resources, being extremely limited, do not come close to revealing the full computational mechanisms that bring the mind about. And furthermore, they are surprisingly limited even relative to our own perception of them. Simplistic heuristics, cognitive strategies, biases, implicit (imperfect) knowledge (etc.) give us a fuller sense of control than we truly possess. Attention determines what we *are* aware of; but what in turn determines *it*, manages it, is itself mostly outside our scope of awareness. All of this concerns even the most seemingly conscious of phenomena. Our linguistic capacity, for example, which appears as high level as possible, does not at all divulge its rules for determining whether a sentence is grammatical or not. Trying to replicate this capacity through programming reveals a richness that is otherwise difficult to notice.

Philosophically, this seeming phenomenological simplicity can be disastrously misleading. It is not that the top, conscious level is *purposely* misleading – quite the opposite. It approximates reality rather *well* (though this is not its purpose but just a means). It sets our attention on what is *important* and that cannot be handled automatically. But the sense behind these "well" and "important" is not at all mathematical nor philosophical; it is a pragmatic, evolutionary sense. In trying to go beyond how things appear to us, to reach how things really are, the hidden computations that construct our perception of reality and our understanding of it can withhold important details. The pragmatic optimization that hides them from us is a philosophical disservice.

The full cognitive mechanism must be taken into account then, at least a priori. What is the place of the conscious in all of this?

> [O]ne of the achievements claimed by contemporary cognitive psychology has been the experimental demonstration that persons have limited introspective access to their own cognitive states and processes… (Greenwood, 1999).

However:

> This is not to say that cognitive psychologists dismissed introspective or verbal reports as nothing more than forms of behavior (Watson, 1925), or as mere sources of cognitive hypotheses having no evidential value… (ibid.).

In modern cognitive science, the relative part of the conscious within the mental is much diminished. However, as part of the mental, it too is rehabilitated from the devastating behavioristic era. The conscious is not only a heuristic guide for cognitive experimental research, but an important, even central, topic of cognitive inquiry. The subjective can be coherently correlated with behavior:

> When… report protocols are implemented judiciously, subjective and objective measures of consciousness are in concordance with each other. (Koch, et al., 2016)

This, in turn, can also be related to the inner workings of the brain:

> The research strategy to identify the neural correlates of consciousness (NCC) involves relating behavioural correlates of consciousness to the neural mechanisms underlying them. (ibid.)

There are many challenges left for relating the subjective to the objective, scientific challenges as well as philosophical ones. But *ultimately, the subjective is also objective*. It is part of what a full scientific, cognitive account must, in principle, fit in. And concretely, the subjective is now also being explored as such, in objective terms.

The conscious/non-conscious distinction, though familiar and by now intuitive, is certainly simplistic, even extremely so. In reality, it is somewhat graded; it is not completely coordinated with the non/attention distinction, nor with objective measures versus subjective reports. And all these may differ with respect to various tasks and operational measures. But the distinction nevertheless *is* an important first-order approximation, central to modern cognitive science.

This concludes our historic-conceptual tour of the cognitive revolution. What implications might it have for the philosophy of mathematics?

1.4 Applying the Revolution

The traditional anti-psychologistic stance excluded the subjective manifestation of mathematics (p. 14) from having any philosophical role for mathematics. As an opening position, this would render the cognitive revolution irrelevant as well. But the revolution itself can be viewed as concerning the very same reasons that brought about anti-psychologism as the reigning dogma in the first place (Chapter 2). And with the logical approach as a form of behaviorism, the motivations and advantages of this dramatic change within psychology may support an analogous shift for the philosophy of mathematics. The form of such a support is not one sharp philosophical line of reasoning. Instead, it results from wide conceptual-empirical scientific progress, which bears a value of a different (and perhaps more substantial) sort (than the intuitionist critique of the logical foundations, for example).

Embedding the philosophy of mathematics within the larger, holistic philosophical-scientific picture could allow it and science to mutually communicate and nourish one another. Analogously to behaviorism, I suggest that logical foundations, too, may have "progressed far enough to recognize that the solution to some of its problems depended crucially on the solution of problems traditionally allocated to other disciplines" (p. 28) – namely mental access to the mathematical realm. In (Hendricks, et al., 2006), for example, many leading philosophers of mathematics expressed hopes for progress to come particularly from our improving understanding of the brain and the mind. A cognitive *approach* aims to head there systematically. And it is the contention of this chapter that this is motivated and indeed *justified* by the cognitive revolution.

With the logical foundations approach considered "behavioristic", let us now figuratively "apply" the cognitive revolution to them, conceptually, to set forth (in broad strokes) a <u>cognitive</u> approach to mathematics and its foundations.

All the fundamentals that were brought in in the previous section bear on the exploration of mathematics. Hopefully, it has been read with that in mind. But this is important enough to warrant spelling out.

Fundamental intuitions cannot be trusted – even mathematical ones. Seeming phenomenological simplicity can hide important details that the mathematician is not aware of. Simplistic heuristics, cognitive strategies, biases, implicit (but partial, e.g. statistical) knowledge, and rich but consciously inaccessible computational mechanisms may all play an unrecognized role. This is a *dangerous* role to begin with, given the presumption of mathematics to transcend evolutionary pragmatics, into the absolute. (For example, when geometric intuition may be grounded in hidden innate mechanisms that are adjusted to life at this scale, might Euclidian geometry thus be falsely perceived as necessary?) Mathematics, perhaps more than any other topic, cannot allow for any content hidden or not explicitly accounted for. This was the motivation for the development of logical rigor, of course; but there may be more to mathematics than that. (All of this applies not only to logical foundations but to phenomenological foundations à la intuitionism too.)

The danger for mathematics goes deeper than just this matter of difference in standard. After all, in the scientific exploration of *physical* reality, cognitive science need not enter. Our mental gateway to *that* reality, though worth exploring for its own sake, is ordinarily *sufficiently reliable* to allow for cognitively naïve exploration as a standard. Evolution, given the importance of proper interaction with physical reality, shaped our cognitive systems accordingly. This reliability *cannot* be assumed where it comes to mathematics at large (higher mathematics included).

The way to assure against the cognitive system misleading us, then, is to go into its inner workings and unravel the cognitive processing of mathematics. This exploration can force us to notice aspects that we have taken for granted. The ultimate assurance is in a complete mechanistic, computational account (including phenomena for which having an underlying computational structure need not be written on the face of them). It is true that mathematics may appear to transcend any computational basis. But as in general, there may be an underlying computational structure to it that, if searched for as such, with hard work, could be exposed. The computational standard – which is seemingly restrictive mainly in comparison with idealized, infinitary set theory – is nonetheless extremely rich and open in potential to new conquests. Striving to adhere to it may not only be of methodological value but may bring us closer to the way things really are (at least in the mathematician's mind).

Mathematics-related "cognitive states and processes", then, "are [or rather should be] as legitimate objects of theoretical explanation as are observable behaviors" (p. 32). To account for a cognitive process (such as counting) in mechanistic terms is not the same as accounting for the *laws about them* that they exhibit (i.e. a purely epistemic foundation). The mathematical laws are taken here as a particular form of behaviorism's "operational definitions". But "[o]perational measures of theoretical constructs … do not in general determine the meaning of such theoretical constructs" (p. 31). Taking ℕ to be "rigidly and exhaustively defined" (p. 26) by its mathematical laws, the logical foundations way, presents their failures in the light of this fundamental deficiency of behaviorism:

Foundations should not merely wish to account for mathematicians' linguistic (idealized and normative) behavior. The mathematician seems to refer successfully to ℕ (if not the set universe) as a distinct structure, and so a foundation should allow for that too (or explain it away, somehow). But this is ultimately a *subjective* expression of mathematics, *ratified* by mathematical practice's actual language, which alludes to that successful reference (as in the definiteness of "*the* natural numbers"). The linguistic level as a logical one captures an important dimension of mathematical reality (as perceived by the human mathematician); it is "the plane of truth" onto which the full mathematical reality, ontology included, is projected. But if this level fails to reflect differences within what they take to be mathematical reality, differences that presumably let them direct their gaze at one intended structure (considered predominantly important) rather than another ("non-standard" one), then for foundation, that level is inadequately thin. The phenomenon of incompleteness only drives the wedge further in, between mathematical ontology as perceived and its logical-linguistic expression. Not only is this complete set of truths insufficient, but it is not available to mathematicians even in principle. Yet the structure itself, in some other sense, apparently is. Where does the conception of truth-value realism come from? How on earth (below the platonic heaven) could it be justified? The subjective conception of mathematical ontology, with the idealizations involved (first and foremost infinity), cannot be grounded in its logical-linguistic expression; but that conception must, by definition, be there, in the cognitive system – somehow.

On a general level: Without going into the cognitive workings, a philosophical theory cannot make room for both the *conscious* level and the grander mental picture of which it is part. Accordingly, non-cognitive approaches in the philosophy of mathematics tend to gravitate toward the two simplistic extremes, either basically accepting mathematicians' own perception of their subject matter or rejecting its relevance[7]:

- Mathematicians essentially know what they are talking about/doing. The foundationalist's job is to articulate the framework in which it all takes place.
- Mathematicians are good at, care about, and may be an authority *in mathematics*, not philosophy – even be it mathematics' philosophy. Theoretical-philosophical considerations alone are of importance.

Cognitive science, instead, lays the groundwork for a more precise consideration of mathematicians' own grasp of mathematics at large and its many particular issues and related intuitions: An account of their mental interaction with the realm of mathematics which *incorporates*, accounts for, how mathematics appears to them – while going beyond this subjective manifestation (if non-conscious *essential* processing is also at play). For a midway slogan, a cognitive approach can:

Take mathematicians' intuitions seriously – but not too seriously.

Within this general narrative of a cognitive approach and the various lessons and implications of the revolution, we should distinguish (following the introduction to this chapter) between two main levels of foundations:

i. A partial foundation that captures (some or all) the mental aspects of mathematics. It may be *instrumental* in achieving specific foundational goal/s and providing *insights*, of various sorts, into mathematics.
ii. A full-blown *cognitive metaphysical* foundation, that grounds all other foundational aspects.

---

[7] See (Benacerraf, et al., 1984 pp. 3-5) for a presentation and short discussion on this distinction between philosophical approaches.

Philosophically, the former is relatively non-presumptuous. One only needs to accept the *potential* relevance of the cognitive to the mathematical (e.g. as a source of insight). Any particular such attempt should then be judged on its own merit. The latter, however, means utilizing the cognitive revolution to provide something deeper than it explicitly offers – an account of mathematics itself as founded *upon* the cognitive (somehow). The next two sections address these in order.

1.5 Representational Foundations

Particular mathematical objects, types of objects, interactions and relations between all of them, and whole structures (with the collections of relations that they bring) are all attended to by the mathematician's mind. Whatever the mathematical entities themselves may be, mathematicians presumably have some sort of mental *representation* of those entities they deal with (according to realism) or at least think they deal with (non-realism). It seems to mathematicians, for our discussed example, that in $\mathbb{N}$ they have a distinct, well-determined structure in mind. This phenomenology, as part of the mental, is grounded in the cognitive system and should be accounted for as such (or at least explained, if they are misled by their non-conscious). This is regardless of whether particular debated (non-standard) numbers or number-theoretic propositions (undecided, e.g. by PA) could ever come to reveal a discrepancy in that perception of definiteness. In general, taking representations seriously suggests, I propose, the possibility of *systematizing* the issue into a <u>*representational foundation*</u>: a foundation that reveals and accounts for the mechanistic representational structure that underlies the mental manifestation of mathematical entities.

Representations of objects (including whole structures) are, by definition, not the objects themselves, and thus this sort of foundation is inherently *meta*-mathematical. The cognitive representation level therefore serves in this way as a foundational basis that grounds the whole of mathematics' *perceived* ontology (metaphysics aside). Like logical foundations (as a particular form of representation), it cannot on a standard realist reading be *object-complete*. There may be concrete mental representations for small numbers, for specific larger ones (e.g. 10^6), and for every one of them, if we accept a sufficiently idealized sense of cognitive representation. And for $\mathbb{N}$ too, somehow. But the problem with extending Frege's numbers into the infinite pertains to any cognitively grounded representation just the same. Although the mind can represent

ℝ, it cannot represent every real number. No cognitively reasonable (meta-theoretic) idealization would suffice.

Representation here is broadly construed, comprising of not only *object* representation but also *knowledge* representation, where furthermore, importantly, *knowledge* is not necessarily just *propositional* knowledge. Thus, a representational foundation also accommodates a *cognitive epistemic foundation* that in relation to the object representations "reveals ~~the true proofs or~~ justifications for mathematical propositions, or provides *one* way in which mathematical propositions can become known" (adapted from p. 10). But the representations themselves need not be logical-linguistic; they are not *grounded* in the objects' propositional properties. Some propositional facts *about* a representation might not be obtainable by the cognitive architecture that suffices for handling the representations. Therefore, there could be a valid representation of ℕ without a *representation* of the truths about it, for example. This is similar to Hintikka's and Shapiro's discussed attitude toward foundations (Section 1.1), where a foundation can be of some value even without a complete deductive support. However, unlike their approach, representation need not by default provide *any propositional* epistemic content. (Would non-linguistic animals be precluded from having any mathematical representation?) These levels are not to be conflated. A representational foundation serves as grounds for the *exploration of the interaction* between objects and the propositions about them, which runs deep. This is an important feature of any cognitive approach.

A representational foundation could conceivably represent all mathematical ontology (at least that which is found in mathematics as practiced), then (as the mathematicians do represent it mentally). However, set theory already does that, does it not? To quote an explicit textbook presentation:

> A typical example of the method we will adopt is the 'identification' of [the geometric line] with the set … of real numbers. … What is the precise meaning of this 'identification'? *Certainly not that points are real numbers*. … What we mean by the 'identification' of [the line] with [the reals] is that the correspondence … gives a **faithful representation** … which allows us to give arithmetic definitions for all the useful geometric notions and to study the mathematical properties of [the line] **as if points were real numbers**. … In the same way, we will discover within the universe

> of sets *faithful representations* of all the mathematical objects we need, and we will study set theory …[9] *as if all mathematical objects were sets*. ((Moschovakis, 1994), emphasis in the original, quote from (Maddy, 2017)).

As such, "identification," "faithful representation," and "as if" remain undefined, left for the reader's intuitive understanding – and necessarily so. As Maddy herself continued:

> The *trick*, in *each case* [emphasis added], is to identify the conditions that a 'faithful representation' must satisfy. For the case of ordered pairs, this is easy: two of them should be equal iff their first elements are equal and their second elements are equal. The case of the natural numbers is more demanding: a set of sets with its operations should satisfy the (full second-order) Peano Postulates. ((Maddy, 2017)).

Identifying the conditions that a "faithful representation" must satisfy is a trick, an ingenuity required per each individual case. Such is the encoding of all of mathematics into the world of sets. It is essential *to* this type of a foundational program yet is not *part* of it. It is unsystematized and necessarily *non-canonical*: As (Benacerraf, 1965) famously pointed out, for the encoding of ℕ into sets, both Zermelo "numbers" and von Neumann "numbers" would do (as would many others). Both {{{}}} and {{}, {{}}} (respectively) *can serve as* the number 2. And hence, being different as sets (e.g. containing a different number of sets), with no justification to rule out one rather than the other, neither of them *is* 2, metaphysically. This is not what the mathematician *means* by "2," which is uniformly fixed across all usages, regardless of extra content that may be to what happens to *serve as* 2. But, to bring this discussion into our cognitive context:

Being 2 and being what is meant by "2" are not necessarily the same. Not being what is meant by '2' is taken as an *indication* that it is not 2. This is a *subjective*, pre-logical requirement. In between the actual metaphysical 2 and {{}, {{}}}-as-2, there is the mental representation of 2. What is meant by "representation" in such discussions as above is, implicitly, a representation of a representation – each one of a wholly different sort; namely, a set-theoretic representation of a mental representation of the actual (platonic) mathematical object. The *definition* of *faithful representation* is with respect to mathematicians' own mental representation; this representation is what it must be *faithful to*. This is all part of what a foundation of mathematics needs to account for. Without this, without a representational foundation, which grounds what mathematics means to mathematicians, things are necessarily underdetermined:

> But how does a set-theoretic ontological foundationalist, for example, go on to *establish*, or even argue for, the claim that *V* is the unique foundation? Since we can interpret every mathematical theory in more than one system, how do we know which is the right one?…
> [M]athematics itself does not decide between the alternative ontological foundations. As far as mathematics is concerned, any of them will do, or we can just eschew an ontological foundation altogether. (Shapiro, 2004)

Interpretations are certainly useful, for the sake of a *meta-mathematical corral*, *risk assessment* and more. However, from a general foundational point of view, they are not transparent. Distancing us from what we mean (e.g. numbers, not sets), they also distance us from the mechanisms that underlie our notions (as well as from the actual metaphysics). The freedom to interpret ontology, to encode it in another, defies uniqueness. This uniqueness is desired because, for the mathematician, it appears indeed to be the case. 2 is a unique object, and mentally, the mathematician can tell the *difference* between it and an interpreted representative. Mathematics itself, in this sense, *does* decide (if we can all speak on its behalf) between the mentally represented object and its interpretation in something else. It is only within a limited frame, with all the concepts fixed and formalized, their faithful representations predetermined, that with respect to proofs, interpretations are fine. But formalization itself can hide details that are mathematically important (as we see in Chapter 7). The details are to be found in the one foundation that is privileged as unique – the mentally "original," uninterpreted, representational foundation.

The point here is not to contribute directly to Benacerraf's problem. The lesson is a general one. Foundational debates and progress interact with and rely on what mathematicians have in mind for guidance. In doing so, they track and gradually reveal the representational foundation (even if this is not framed in such terms). This includes, for example, categorical and univalent foundations, which are arising as an alternative foundation to set theory and reforming the whole foundational debate (bringing back into it mathematicians who did not relate to its strong logic + set theory disposition before). This type of foundation is also susceptible to, and would ideally converge with, a representational foundation. Anything that mathematicians take to make sense mathematically thereby makes sense in their minds. This is the *ultimate universality* of a representational foundation. The representational foundation is something that *essential guidance* (p. 10), a central advantage of category theory, may interact with and even be grounded in. Whereas for Maddy:

> It seems that both these camps [the set theoretic and the category theoretic] are chasing a false goal: a foundation that delivers Essential Guidance, a single understanding of what mathematics is, a single recommendation on how mathematicians should think. (2017)

"Single" sounds restrictive. But the breadth of what computational cognition can support is no more restricted than all that human mathematicians could ever achieve. To unify it all under its mental representation is not to restrict it. This is the arena in which to explore the *many* ways mathematicians think and could think and the interactions between those ways – the interaction between category theory and set theory (as foundations) included.

As the narrative developed in this chapter suggests, all this is not to break away from the logic-based tradition. It is to take it further, deeper into self-reflection beyond the logical-linguistic level, which could further serve the exploration of mathematics. That traditional logical basis – though it did not stand up to the foundational expectations placed upon it – is still central and useful for the exploration of mathematics. Likewise, it might be so for the representational basis of mathematics. Even if in general the metaphysical and epistemological foundations of mathematics are not reconceived into any cognitive terms (but are kept, e.g., platonic and purely-logical respectively), there is still room for achieving various partial foundational goals and providing insight through cognition. The supporting computational structure, even if it does not determine ontology and metaphysics, is at least a central gateway to their exploration.

Exploration through representation is the main substance that this dissertation offers in Part II. But now, let us move on to how the cognitive revolution may suggest a different *metaphysical* foundation for mathematics.

1.6 Cognitive Foundations

From our broadly construed foundational point of view, the cognitive revolution calls for a re-evaluation. Logic as a simple and elegant foundational basis may suffice when one does not set out to make the full manifestation of mathematics relatable to it, the subjective included. But in this chapter, we roll back on what *foundations* has *come* to mean (while absorbing their limitations) and lift any restrictions on its ambition that it may have come to have. In this section in particular, let us reconsider the grand enterprise of a metaphysical foundation (to which all aspects and manifestations of mathematics ultimately relate).

Logic had to be developed as a formal tool so that it would be rich enough to capture an important aspect of the manifestation of mathematics – mathematical language and its structure of truth. In those days, it was the case that "Only mathematical language seems sufficiently general, intersubjective, and in fact objective, to be a candidate for reflecting the essentials of the practice and, through it, grounding all of mathematics" (p. 16). Mental mathematical content was only a subjective manifestation (p. 14). Today, however, the subjective – grounded in the objective (the material brain) – is (beginning to be) explored objectively in scientific terms. And computation is not only assumed to be at its core; it is actually turning out a rich-enough tool suited for animating the mental realm and all its riches. This, in cognitive science, promotes a metaphysical *parsimony*:

As a general working philosophy, cognitive science is driven by *naturalism*, taking all that is not physical to be mental, in turn reducible to the physical (at least in the functionalist sense of abstract mechanisms being realized by physical reality). Beyond these two realms (and the riddled connection between them), this working philosophy may not, on philosophical principle, rule out the possibility of other, distinct metaphysical realms, e.g. a moral *realist* landscape. But if such a realm is to be reflected and represented in any way in the mind – if living beings are to interact with it mentally – then cognitive science would be the field to explore this. And such a presumed realm, neither physical nor mental, does not offer science any grip; it is outside its purview.

This renders that supposed realm scientifically dubious. For mathematics, this is the problem behind Benacerraf's challenge:[8]

> [T]o provide an account of the mechanisms that explain how our beliefs about these remote entities can so well reflect the facts about them. (Field, 1989)

This standard cognitive science stance, applied to mathematics, would then suggest that *there is no such remote, platonic realm*. And indeed, this position is a common one among cognitive scientists (I dare say) and particularly among the very few who are taking on the philosophy of mathematics (Chapter 3). It seems to collide, of course, with mathematics' own working philosophy, as well as with the mainstream philosophy of mathematics.[9] Resolving this conflict is *the main task of any cognitively informed philosophy of mathematics,* which cannot simply leave Benacerraf's challenge for others to resolve, somehow. Any approach that naively relies on what mathematicians as humans see mathematics to be – but does not offer, in principle and in practice, a way for cognitive science to go deeper – is inherently flawed.

Within this naturalistic working philosophy, the options for grounding mathematics are in:
(P) The physical.
(M) The mental.
(P-M) The interaction between the physical and the mental.

From a general scientific point of view, (P) is quite natural. Cognition aside, mathematics seems to be integrated into the physical world itself (I attend to this integration in Section 2.8). Presumably, we have the ability to pick up on mathematics, then, in the same way as we pick up classically physical information from our environment. But as discussed (p. 14), in the physical, mathematics is expressed *implicitly* and so cannot serve as a foundational theory. Mathematics' boundaries must be drawn first, explicitly. Only then could we hope to reverse engineer it cognitively, perhaps, to figure out how implicit mathematics could in *some* sense give rise to explicit

---

[8] See also Benacerraf and Putnam on "The problem of 'access'" (1984 pp. 30-33)
[9] For naturalism in the philosophy of mathematics, usually without much concrete relation to cognition, see e.g. (Maddy, 2005; Paseau, 2016).

mathematics. (If that sense is to be taken as metaphysical rather than just developmental or computational, then the starting explicit mathematic manifestation serves as a foundational ladder that we then decree as metaphysically inessential). Thus, (P) requires some preliminary cognitive, non-metaphysical foundation for a theory. But this is a delicate point, and the term *cognitive* should be avoided with respect to (P) – which I do not attend to further. Still, (P) too opens the question of how we come to perceive the mathematical aspect of reality. As a cognitive science exploration, it would be very different from the standard exploration, of the cognitive interaction with (classically) physical reality.

As for both (M) and (P-M), they are inherently cognitive to begin with, the latter fitting *embedded cognition* (Wilson, et al., 2017). In terms of parsimony, (M) may be the more natural option. Mathematicians do have mental entities anyway, which correspond to specific numbers, to ℕ as a whole, etc. And these seem to *apply* to non-physical, mental entities just as well (e.g. counting imaginary sheep). Why bring in the physical at all, then? First, perhaps in order to explain the place of mathematics in the physical sciences. Another rationale might be that our general learning skills, cognitively supported by our nature, may be inherently outward looking, empirically based. In any case, for both (M) and (P-M) (though not (P)), it suggests the possibility and perhaps the necessity of the following:

Cognitive foundation: a foundation that *grounds* mathematics in cognition, metaphysically; grounds mathematical ontology itself in the mental entities and supporting mechanisms (as with respect to the world, for (P-M)). The natural approach is to rely on a *representational* foundation, which already systemizes the mental manifestation of mathematics. That basis is then leveraged to a metaphysical status. The mental "manifestation" of mathematics is a *constitutive* representation of the perceived ontology. In principle, this move alone could still keep to mathematical ontology as standardly conceived by mathematicians. It would be a meta-mathematical grounding, not *object-complete* (p. 18), that posits a further metaphysical gap between the cognitive ontology (which determines and indeed constitutes P(ℕ), say) and the mathematical one (which contains every subset of ℕ). However:

Cognitive *naturalism* as applied to mathematics implies furthermore a purely *cognitive metaphysics*. To spell this out: There is no further metaphysics beyond the cognitive.

The cognitive foundation is object-complete; mathematical yet non-mental objects do not, strictly speaking, exist – not in the "true" foundational metatheory, but only in the mathematician's theory. This *gap* between the cognitive ontology and the perceived or presumed standard mathematical ontology can no longer be left an unaccounted-for metaphysical one. It must be reduced to a *cognitive* gap: mathematical ontology is to be brought about and *bridged* to by the inner workings of the mind itself.

The philosophical ramifications are substantial. But this does not mean a straight denial of the foundational relevance of mathematicians' commonplace fundamental intuitions. The realistic talk – which refers to an autonomous reality, relative to which the sentences uttered are rendered true or false (proofs serving as mere witnesses)[10] – should still be accounted for, ultimately. The particular cognitive foundation, and its implications for the generated mathematical reality and for mathematical practice, then depend on the details of that bridge and what we make of it. A detailed cognitive account of the bridge *may*, first, teach us that modern mathematics has generally got it right (despite not accounting for the cognition that underlies that top, realistically perceived level). Or, second, it may reveal limited biases and misperceptions and, through this analysis, lead to a refined and improved account of mathematical reality. Or, lastly, it might call for a drastic revision of the structure of mathematical reality and, with it, of mathematical practice.

In any case: Mathematics has come a long way since the time it was almost synonymous with calculations and concrete geometry. But so has computation and what the mind can be seen to achieve upon it. All modern mathematics – this naturalism would suggest – can still be reduced to that. Russell's philosophical claim that mathematics can be reduced to logic necessitated a large body of technical, non-philosophical work to support it. The same holds, possibly to a much larger degree, for the philosophical progression of claims, from "mathematics is systematically *represented* in the mind" to "mathematics is *grounded* in the mind," and last, to "mathematics can be *reduced* to the mind". The purpose of this dissertation is not to complete this journey but just to

[10] When it comes to mathematics, just as in cognitive science at large: Language's relation to non-linguistic cognition and to objective reality should be cognitively accounted for. At least as a default assumption, it should be theoretically grasped uniformly over mathematics and the rest; otherwise, the differences must also be cognitively accounted for.

establish and embark upon it, focusing on the first stage – the quest for a representational foundation, with the limited, yet real, profits this might already offer.

The first step on that long journey would be to dispel the rather sweeping objection to approaching the foundation of mathematics in *any* cognitive terms. Let us now attempt this step while developing the themes of this chapter along the way.

2 Beyond Anti-Psychologism?

Cognitive science seemingly gives us empirical knowledge, of a contingent nature. On the other hand, the standard conception of the mathematical universe is that of a timeless, abstract, non-human, non-contingent, a-priori explorable platonic one. Thus, approaching mathematics (as opposed to mathematical thinking) through cognitive science might be considered hopeless from the start. This may be the view of (at least) analytical philosophers as it is true for mathematicians.

On the philosophical front, Frege and Husserl's famous attacks on psychologism "for a while… even rose to the status of a paradigmatic achievement in philosophy" (Kusch, 2015). However, within the relevant literature, "this assessment of antipsychologism is now widely contested"; "many recent re-evaluations… conclude that Frege's and Husserl's arguments are question-begging". I re-evaluate the old arguments against psychologism in the face of modern cognitive science (where leaning toward psychologism is much more common). Through this narrative, I further develop the first chapter's general theme of a cognitive approach to mathematics. Because I aim to keep the discussion within the context of the present project, I keep the engagement with modern literature to a minimum, addressing the issues only through the introductory (Kusch, 2015). All quotations are from there.

Among mathematicians, anti-psychologism probably reigns as strongly as ever. One might think this prevalence is irrelevant in assessing the philosophical arguments against psychologism and hence to the positions and their viability. But my approach, recall (Section 1.3.3), is different. The mathematicians' intuitions are part of the data and need to be explained (while stressed as just a part of the story). Why is it that they find anti-psychologism so appealing? Its wide acceptance goes beyond the direct influence of the philosophical world (I dare say).

Much of the original anti-psychologistic debate has been about logic and was meant (for Frege) to apply to mathematics also through its grounding in logic. But we, in this dissertation, are concerned directly with mathematics and certainly do not abide by some form of logicism. Thus, I recast all arguments as being about mathematics itself

– again, historical scholarship aside. What matters to us is capturing the reasons why many philosophers and most mathematicians still find anti-psychologism *regarding mathematics itself* appealing. Of course, *behind* these scholars' reasons, there is a cognitive system that may be leading them astray. Their reasons may thus be *explained away* (as with Lakoff & Núñez, Section 3.3). Cognitive science may call for such "cynicism" toward their perception of the mathematical realm. However, the challenge I undertake here is quite the opposite – to respect this orthodoxy as much as is possible (in the face of modern cognitive science). I argue for the metaphysical *possibility* for a *cognitive foundation* to stand at the base of a *realistically conceived mathematical ontology* that manages not to conflict with mathematicians' basic conception of their own subject matter (even if going beyond its phenomenological limits and enriching it).

Anti-psychologism was a perfectly reasonable position at the turn of the previous century given the scientific state of psychology at the time. In this chapter, I show how modern cognitive science can help ease the classical, justified psychologistic concerns. I review those concerns (advancing, roughly, from the simpler to the deeper ones) and briefly suggest why they no longer must rule out a cognitive foundation of mathematics (albeit necessitating a high level of metaphysical commitment, from Section 2.5 on). The actual substantiation of such a promise would require deeper work (than this general kind of philosophical argumentation), the kind of work that is only *begun* by this dissertation. The current chapter concludes with one central threat (to any such cognitive foundation) that remains standing, requiring deep metaphysical work far beyond the project started here.

2.1 Mathematics as Mediated Through Cognition

A central consideration for Frege, which might still bring some to dismiss anything but complete anti-psychologism as a non-starter, is the following, first Anti-Psychologistic concern:

> AP1) Whereas mathematics is the most exact of all sciences, psychology is imprecise and vague. (Kusch, 2015)

The main anti-psychologistic concern in this chapter is metaphysical (regarding whether the mathematical realm somehow is a psychological one), and we will soon get to that. But let us here take the AP1 phrasing verbatim, epistemically, as being about the *scientific fields*, fields of inquiry with their immensely different methodologies for truths finding. AP1, as a statement of fact, is certainly still true in general (even if psychology has greatly progressed since the days of Frege). But such generalizations are of no use here. The bigger project of scientific knowledge is holistic. If knowledge from a generally less precise field can nonetheless teach us something about a generally more precise one, we should take it. To take the precision-ordered state of affairs between the sciences as a prescription (regarding which science could support which), rather than mere (generally valid) description, would be dogmatic. As long as mathematics, on its own, is merely *the most* exact but perhaps still not *perfectly* exact (in some sense), it should take any help it can get.

This goes beyond just precision. That empirical, scientific findings (in the cognitive realm) may tell us anything at all about non-contingent matters (in the mathematical one) might perhaps not be so obvious – and, for the most part, is not practiced as such. Yet this is part of a much wider set of phenomena, of facts that are deeper, 'more-necessary' than the methodology that finds them. A computer may prove a mathematical result – which is not to be trusted more than we trust the validity of the program, the operating-system, etc. – while the result itself, if true, is as a priori and non-contingent as any other. Mathematical facts can also be noticed through sheer observation, on simulations or on naturally occurring phenomena[11]. And evolution, too, converges on (abstract) engineering principles (i.e. how to build a functioning eye) or even epistemological ones (the brain does quite a lot of that) that are now explorable in combination with biological methods – A.I. drawing ideas from neuroscience (Hassabis, et al.).

Even with no metaphysical commitment (of any sort) to mathematics *as* psychological, research into psychology may then still, in principle, be relevant to the exploration of

[11] "While examining friendship between children some fifty years ago, the Hungarian sociologist Sandor Szalai observed that among any group of about twenty children he checked he could always find four children any two of whom were friends, or else four children no two of whom were friends." (Gowers, 2008 p. 562)

mathematics. [12] It is particularly relevant as our gateway to the mathematical realm (and being a gateway is perhaps what makes it a contender for being what mathematics is).

The whole of Chapter 1 was about this. So just to bring it together with the classical psychologistic debate:

Much within all that is mathematical (p. 14) concerns psychology (e.g. visualizations or intuitions). However, even for its foundational core, even when confined into logic, our knowledge of it is still psychology-dependent. Even leaving aside (as I do) the big topic of psychologism regarding logic itself, there are still matters of formalization and the ("apodictic") *self-evident* status of axioms. Self-evidence would seem to be a feeling, a psychological phenomenon. As Husserl notes (Kusch), this does not entail that the involved truth itself is psychological, just its perception. And Frege, in his anti-psychologism, takes self-evidence as something distinct from mere obviousness (Snyder, et al., 2018 pp. 101-102). But what other access to the realm do mathematicians have? Can psychologistic self-evidence be avoided? That anti-psychologism's ambivalent reliance on self-evidence can be question begging or even contradictory is a classic criticism (see e.g. Wilhelm Wundt's or Moritz Schlick's, in (Kusch)). But back then, self-evidence seemed like a necessary starting point ("If we were not allowed to trust self-evidence any more, how could we make, and reasonably defend, any assertions at all?" (Husserl, in (Kusch))). Modern cognitive science, however, puts us in a very different position (Section 1.3.3). Naïve reliance upon self-evidence is thus no longer reasonable – nor *necessary*. Cognitive science could thus conceivably shed some light on mathematics. (AP1), though still true, is irrelevant.

2.2 Mathematical Foundations for Cognition

Moving on to psychologistic metaphysics, a familiar worry is the following:

> AP2) If mathematical rules were based upon psychological laws, then all mathematical rules would have to be as vague as the underlying psychological laws. ((Kusch), adapted from Husserl)

---

[12] In the end, science, too, is based on mathematics, and so we might not be able to avoid the inherent circularity – but we *can* clear some misperceptions by going deeper.

Or, to focus on a simplified variant (avoiding the extra metaphysics of "based on… underlying…"):

> AP2') If mathematical rules were psychological laws, then they would have to be as vague as psychological laws.

Technically, as philosophical arguments, these are question begging. As J.J. Katz says: "One might well reply that … logical and mathematical laws are the psychological laws that are exceptions to what is otherwise the rule in psychology for now" (in (Kusch)). Or as Moritz Schlick states: "One sees immediately that one might with equal right infer the opposite: since logical structures, inferences, judgments and concepts undoubtedly result from psychological processes, we are entitled to infer from the existence of logical rules that there are perfectly exact psychological laws as well" (in (Kusch)). Still, this worry admittedly *does have* some appeal and is part of why anti-psychologism is, to some extent, a standard view among mathematicians – philosophers' debates aside. The cognitive revolution at least now offers something more:

Long gone are the days of behaviorism in science. Mental states are part of theoretical deliberations and modeling, and neurons and their aggregates stand at the core of the vast and exponentially growing field of brain science (now almost interchangeable with "neuroscience"). Psychology, though still a (so-called) soft science, is now underlain by a hard science. Current mathematical models of the brain may still be just approximations (which could be overthrown by discoveries of other biological/chemical/physical influences that are relevant to the computational behavior of the system). But its foundation is grounded in elementary physics and is ideally mathematical. And taking a functionalist informational view, the mathematical modeling is all that matters (deeming irrelevant the details that are outside the (ideal) mathematical model). Psychology overall might not be completely reducible to such a *mathematical foundation*, even in principle. Some, or even most, of the mental realm may concern emergent properties that are doomed to stay imprecise and vague forever (even if not as much as they seemed over a century ago). But at the very least, a mathematical foundation does *make room*, in principle, for this other type of phenomena – mental, yet perfectly mathematically precise.

What a mathematical foundation of cognition itself now gives us is the conceivability of a cognitive approach to mathematics through circularly coherent holism. Cognition has a mathematical core, which could (perhaps, somehow) serve as foundation for mathematics as we know it – which in turn allows us (among many other things) to scientifically account for that core (the "elementary particles" of cognition). No longer must vagueness and imprecision, classic central features of psychology (or rather "bugs"), stand in the way of such an approach.

This is not sufficient for dispelling the appeal of AP2' completely. Whether part of psychology or not, mathematics is very different from any other psychological realm. Without an account of *how* mathematics is part of psychology, including an account of *why it is* very different from everything else psychological, such debates are fated to remain vacuous. But conceivability is a necessary preliminary. Let us push this line further on our way toward (the possibility of) a cognitive foundation.

2.3 Ontological Rigor

To pick up on the first chapter's theme of a mechanistic standard of explanation: Mathematics, whether 'psychologic' in any sense or not, needs to be based upon something as precise or more (if at all). But we already have that in logic. Rigor through logic is what is supposed to make sure we had everything, all mathematically relevant details, laid out. This picture rests on the view that whatever is foundationally important about mathematics is part of our conscious perception of mathematics – or at least sub-conscious (i.e. can be "dug out" by the progress of the practice). Anything below that – though it may certainly govern intuitions, visualizations, etc. – concerns only the psychology of mathematics but not its (metaphysical) foundations; not (strictly speaking) mathematics itself. To be sure, a great deal of math is occurring in the brain, but that is the system doing mathematics – not the person, consciously, not what we mean by *someone* doing mathematics (I elaborate on this in Section 8.2.1). These two levels of "mathematics being done" are taken to be two separate issues (even though some of the brain's mathematics is what brings about the person's). This conscious-grounded view is a cornerstone not only of logicism/formalism, but of intuitionism too. Even if, for the intuitionist, logic only codifies and reflects what is really at the core – which are the mental constructs – these are also just the phenomenological ones, consciously accessible. That is as deep as one needs to go for foundations.

The question is: is limiting foundations to what mathematicians can "see" going on with their mind's eye, in their consciousness, justified? Are the non-conscious inner workings simply a categorically different, unrelated realm? Or could substantial cognitive operations be non-conscious, hiding from them some of the mathematical complexities involved? In cognitive science, too, "Integrating and manipulating abstract units of meaning (e.g., numbers, words, objects) is a set of cognitive functions that is widely considered to require consciousness" (Hassin, et al., 2014). But the full empirical picture is very much unexplored territory, and evidence is starting to accumulate at least that consciousness is not *required*. For example, Sklar et al. (2012) show that we can "do" two-stage subtraction problems (e.g. 9 – 3 – 2) while not being aware that anything is going on but being primed to name the correct result more quickly than an incorrect one. (Being multi-staged rules out the possibility that this is just implicit *memory*). It seems reasonable to consider this a case where the person herself does abstract, symbolic arithmetic par excellence (not just some sub-system running some concrete calculation). Following this study (and another on sentence-reading), (Hassin, et al.) concluded that "[W]hile early evidence might have suggested that consciousness is necessary for integrating abstract units of meaning, recent evidence seem to challenge this view significantly by providing evidence for the integration of numbers, words, and visual objects". Such findings show, at the very least, that there is a substantial *interaction* between the mathematics we consciously perceive and non-conscious analogs – consciously gained abilities transferred to the non-conscious, made phenomenologically invisible. But the role of such transference is not just to reflect or serve as backup. Instead, it is to free consciousness of such burdens on its load. The question is therefore reinforced: might some (or even many) of the details of mathematics be hidden from consciousness?

This consideration forces us to dig beyond the introspective standard of the aforementioned classic foundational approaches and into the cognitive inner workings. What is required is a standard for when all (potentially) relevant details are accounted for "all the way down" – for when sufficient cognitive depth has been reached (then taking it for the basis). Debates regarding standards of explanations in cognitive science abound. But we need not commit to much here (though I do take it to be mechanistic, in accord with common standard in cognitive science). What is important is that this basis be mathematical (as discussed in the previous section). For example, a mathematically specified neural basis may do (unless advances in brain science tell us differently). This could be a particular form of neural network or (more likely) some higher-level, more abstract computational characterization of such networks. (Deep work would be required to give a proper account, but I am just positing the ideal here).

This level of account (for mathematical objects and structures as manifested/perceived cognitively), I put forth as the notion of *ontological rigor*. Roughly speaking, giving such an account of objects and structures means that we fully understand what we are talking about (and how it relates to other things), in a sense that the familiar definitional standard need not necessarily capture (as I demonstrate through the case studies in Part II). Conversely, a difficulty in giving such an account, no obvious (possibly) underlying cognitive processes having been made explicit, means that the object or structure ought to be better understood, and perhaps reinterpreted. To be sure, mathematics not only currently does but *can* take place with some of the computational details hidden, unaccounted for. Mathematicians can still *use* and rely on a piece of cognitive machinery as a sealed module, a black box whose implementation is opaque. This is analogous to how mathematics faired quite well (to say the least) before, without the modern standard of logical rigor. Perhaps what was being done was not legitimate by modern standards – but it was *legitimizable*. The notion of ontological rigor is a direct continuation of this, driven by the cognitive critique of the logical standard's insufficiency (for achieving what it has aimed to achieve). However, going deeper than what is only conscious/sub-conscious means that the ultimate account no longer must seem to the mathematician to capture what she really has in mind, to *synchronize* with her phenomenology. It only must *incorporate* that phenomenology, account in particular for how that comes to be her phenomenology.

How mathematicians (as humans) come to successfully use, understand, and reason about such non-rigorized “black-box’ components *is* related to their actual cognitive machinery (e.g. controlled, conscious processes turning automatic), and it *should be* explored as such. However, this is integrated with the drive toward ontological rigor, raising the grand technical challenge of basing the whole of mathematics upon a computational basis. Once the full mechanisms are laid bare, unveiled up to their mathematical bone – this puts us back into the realm of mathematical precision and, more generally, will bring us back to *realism* (Section 2.6). With such a standard for cognitive accounts of mathematical ontology in place, we can now re-visit the anti-psychologistic worries and attempt to dismantle them.

2.4 Towards Objectivity

Much in psychology concerns the personal, the subject's subjective view. Not only in terms of the first/third-person distinction but in the variability, for different subjects, in their perception of some subject matter (etc.). And when the subject matter is mathematics, this variation would be a crucial obstacle to its alleged objectivity if psychologistically perceived. Hence the following two possible objections, originally by Frege and Husserl respectively (putting aside Frege's own quarrel with what "ideas" are):

> AP3) Numbers are not ideas since they are the same for all subjects. (Kusch)
>
> AP4) All forms of psychologism are forms of "species relativism". (Kusch)

These rest on the assumptions that *all* that is psychology-grounded varies among subjects or species. This assumption, however, is unwarranted.

Idealized computations, to begin with, are perfectly objective (impersonal) and mathematical, whether introduced and considered abstractly or as abstractions from the workings of a living brain or a computer. Differences in computational structures and in the values they operate on (either their own "innate" parameters or as acquired from the environment) – that is where subjective variety comes from. But not everything varies. Some psychological facts or constructs are intersubjective (shared by many) to varying degrees:

- They may be sweepingly intersubjective, concerning cognition itself, in its definite computational generality (e.g. no appropriately-limited cognizing being could solve the general halting problem).
- They may govern all neuron-based cognition (e.g. bounds on thought-complexity due to energetic constraints, or at least tradeoffs).
- They conform to the constraints of evolution, which brings in a normative component. (For a simplistic illustration: Might some species come to think that 2+2=5?)

- Other commonalities can concern human cognition. A case in point here, separating (AP3) from (AP4), is the possibility of a universal grammar. Such a construct *would* be the same for all (human) subjects – and yet be perfectly cognitive.
- Other facts of cognition may concern only humans with the required abilities and acquired experience and education for it to be relevant to them.

Telling such possibilities (and others) apart might be impossible in practice. Or it may require hard scientific (including empirical) work (consider the debate on the actuality of a universal grammar). But the relevant anti-psychologistic arguments can hopefully be seen not to be necessary.

This brings us to our next theme. Mathematics, if concerned with the fundamentals of cognition, could in principle be cognitive yet still the same for every cognizing creature. It can be *non-contingent* – not depending on the evolutionary path, the individual's upbringing, and general environment nor the random noise that is part of life (including the functioning of any living brain).

2.5 Mathematics as About Cognition

Platonism posits the mathematical realm as completely independent of the mind, and a *cognitive* foundation rejects that. But there are different senses of how mathematics *could be* related to the mind. A central distinction here (stressed by Posy (Forthcoming, 2019)) is between what underlies, or even governs, some area of discourse, and the content or subject matter of that discourse (what it is "about"). For mathematics, the former sense – constraining mathematics and its objects to the mentally representable (e.g. object-completeness, Chapter 1) – has been central to the familiar cognitively motivated approaches (e.g. constructivism, Section 3.1). In this sense, Husserl's following objection (with "mathematics" replacing "logic") may be considered a categorical confusion:

> AP5) If mathematical laws were psychological laws, they would refer to psychological entities. (Kusch)

As Meiland explains, "Just as the laws of planetary motion are not about gravity but are true in virtue of gravity, so too… could claim that the laws of [mathematics] are not about the mind or about standard psychological properties or of the nature of the mind". (In (Kusch))

But instead, turning to the latter sense, let us consider biting the bullet. The standard accusation against anti-psychologism, of question begging, applies to AP5 too. Why couldn't mathematical entities, the ones referred to by mathematical laws, truly *be* mind-related entities, if of a certain and unique kind? They may be mind related in a "Kantian"[13] sense of reflecting certain aspects of cognition itself and cognitive activities (e.g. results of counting). According to such a sense,

> MAC) Mathematics is *about* cognition. It is not directly about the world – but about the computational forms of *representing* it, *through* cognition.

Cognition here is interpreted in the most general computational sense. A weaker, physicalist alternative takes cognition concretely, biologically, as it contingently happens to be.

> BMAC) Mathematics is not directly about the world – but about the forms of representing it, through *actual* cognition, as brought about by evolution.

This biologized version bites the bullet of "species relativism." The contingent cognitive makeup now very much matters. But it is important to stress that:

- Evolution itself can bring in a normative component. It is not just about what the mental structure *happens* to dictate, but possibly about what it is *shaped* to dictate. Actual cognition comes to *approximate* deeper laws. It might not be able to conceive of 2+2 as anything *but* 4. Possibly, because this is true of physical objects. Or perhaps because it is true of the representational computational operations, that are too fundamental for life to be able to be misrepresented.
- BMAC, like MAC, is about the metaphysics of mathematics, not its epistemology (not directly). Mathematical laws that govern cognition need not apply only to those beings with the ability to gain explicit *knowledge and understanding* of it (i.e. humans only, allegedly).

MAC is a metaphysical commitment that goes far beyond the rather scientifically accepted point of view presented in Chapter 1. Just clarifying what it could mean, and

---

[13] To align the terminology with the popular conception of Kant's philosophy of mathematics, leaving aside his actual views and Kant-scholarship. See (Posy, Forthcoming, 2019).

certainly substantiating such a view, would be a great feat, which I hardly attempt to initiate in this dissertation (of which the main aim is not metaphysical anyway). But for the rest of this chapter, let us just see how it could fare with the remaining anti-psychologistic objections (deferring the main motivation for it to Section 2.7).

> AP6) The laws of mathematics do not have any factual content or “empirical extension”. (Husserl, in (Kusch))
>
> AP6') Mathematical truths are not empirical truths. (Frege, in (Kusch))

(These are not necessarily the same. AP6’ could well be about methodology. But here the focus is on metaphysics, so let us treat them together).

A reasonable and modernly relevant reading of this, I suggest, is regarding "empirical" as the *inputs* of the cognitive system, what it perceives about the physical world *outside* it. Much or even most of the brain might indeed be dedicated to the interaction with the individual's body and environment. However, this is certainly not an obligatory property (and the cognitive revolution (Section 1.3) undermines such radical empiricism). The mental can turn inwards; meta-cognitive perceptions and management abound. Previously addressed anti-psychologistic arguments concerned aspects of mathematics for which it was indeed unique, and so they seemed reasonable even if formally question begging. Here, however, mathematics is not alone. A perception of one's own emotions may not be an empirical matter (even if it happens to also have an empirical extension, e.g. in the form of tears). And we certainly have some control over our own thought process, not just our body.

Mathematics, then, could conform to AP6/AP6' (along the tradition of Rationalism) by being a meta-cognitive, "inner mind thing" (to stress how much more must be said). However, meta-cognition can be tricky, bringing in the empirical and contingent through the way the mind happens to be (i.e., despite not depending on system inputs). A meta-cognitive law might still carry *implicit* factual substance (AP6); a meta-cognitive fact might still indirectly be true by virtue of empirical facts and thus, in a sense, be an empirical truth (AP6’). This is indeed the case for all non-mathematical meta-cognition, but would pose a problem for the standard conceptions of mathematics. Where empirical, contingent details can hide is in the workings of the unaccounted-for black boxes. Ontological rigor is thus what could (conceivably) take the dependence on worldly contingencies out of the equation (so to speak). Without this, basing our mathematics and certainly our foundations on introspection alone, we can never be sure. This brings us back to our discussion on apodictic self-evidence and to an anti-psychologistic objection of Husserl, along similar lines to AP1 and AP2[14]:

> AP7) If the laws of mathematics were psychological laws, then they could not be known a priori. They would be more or less probable rather than valid, and justified only by reference to experience. (Kusch)

From our modern perspective, we can no longer simply accept the a priori as a given, trustworthy basis we could build on. To keep things simple, consider the definition of the a priori as "knowledge or justification independent of experience or empirical evidence"[15]. As discussed, we must break the latter two apart. (Direct) empirical evidence means sensory input for the system, from its environment. Experience, however, is not only experience with the outside world – but also experience with one's own cognitive system. And, also as discussed, some cognitive laws could concern the computational basis of cognition (e.g. counting?) and therefore be perfectly mathematical: universal to cognition and hence absolute and necessary. Conceivably, these could thus be experienced and perceived as such (somehow), taken as absolute. In fact, taking for granted (usually implicitly) the true-so-far (until it turns out that it should not have been) may be a rather fundamental cognitive heuristic. To be sure, this might also be done falsely, for laws that only appear to be absolute. Or this may falsely be done for laws that are in fact subjective, specific to the individual, or its species. The question could just as validly be not how we come to see some things as absolute, but as *not* absolute. Subjective differences are filtered through the mathematical community and history. How mathematicians together ultimately came to detect true universality over just species-intersubjectivity (assuming they in fact succeed in this) is a deep matter, which we need not account for here. This is necessarily a problem for naïve reliance on a-priori knowledge too (and so does not support the AP7 objection).

---

[14] Now comparing mathematics not against psychology but against everything else, physics included.
[15] Adapted from (Wiki-APiori, 2018).
The notion itself is not without controversy and has been central in philosophical debates from Kant on. But here I avoid the vast philosophical background as usual and focus only its relevance to anti-psychologism from the modern cognitive point of view

Furthermore, if there is any hope of telling apart the biologically contingent from the cognitively necessary, it is precisely through the kind of inter-species explorations and comparisons, and systematic accounting of cognitive possibilities, that cognitive science is in charge of. What is important here is that the laws of mathematics *could* be psychological, yet valid, and known a priori – in the sense of being dependent on experience that is in fact only with the cognitive system, not with the outside world.

Much still needs to be accounted for if we are to make sense of the a priori. Crucially, this includes the fundamental tool for tracking mathematical reality – *proofs* (how they relate such cognitive-mathematical laws to other ones and how this is perceived as absolute). My modest point here is merely to stress how the likes of AP7 – rather than drive a wedge between mathematics and psychology – should in fact drive the foundational exploration toward the interaction between them.

MAC is brought in only as a natural way (in my view) to withstand the anti-psychologistic objections. Mathematics being a meta-cognitive matter furthermore allows to bring it together nicely with its *abstract* nature (rather than just take that nature as a given). Being about the form of the cognitive process is the way to abstract from whatever its content happens to be. Being about the ordering of objects, for example, abstracts from what these objects truly are. There is no need to go further with it here, however, as mathematics' abstract nature could hardly be used as an argument against it being psychological. As a continuation of the meta-cognitive discussion above, the abstract, too, is common to other areas of the mental realm besides mathematics.

If the conceivability of a cognitive grounding of mathematics along such lines is granted, we are now ready to address the conceivability of a common, central requirement of mathematics' metaphysics – realism.

2.6 Cognitive Realism

We now reach what is perhaps, for mathematicians, the central argument that can make them swear by anti-psychologism before the debate even starts:

> AP8) Mathematics is discovered – not created. Thus, it is independent of us.

As Frege emphasizes, (in the words of (Kusch)):[16]

> [C]oming-to-know is an activity that *grasps* rather than *creates* objects. This choice of terminology is meant to bring out that *what* we come to know is (usually) independent of us. After all, when we grasp a physical object like a pencil, the object is independent both of the act of grasping and of the human actor of the grasping.

The two key issues here are with *created* and with *independent of us*.

Creating as an activity that originates in the mind (and from there, possibly into the world), producing something that does not have to exist, can have a firm intuitive grip on us. It is, however, quite a challenge to make sense of creativity upon a fully specified computational framework (which in a sense aims to take creativity out of the equation). Hopefully there is room for creativity nonetheless. But what is important here is that creativity is not a *necessary feature* of *all* that is psychological. As I have suggested, there are *laws to cognition*, both to particular biological systems BMAC and to cognition at large MAC. Such laws are not in any sense created by us, not dependent on our "will" or creativity. And in the case of MAC, they are not only well determined, but fixed for all eternity, even if no creature ever has lived. In this important sense, they are indeed independent of us. "Will" and creativity may be needed for the practice, for exploring this realm and *discovering* its laws, making them explicit and unveiling reasons for their being so (attesting to their lawfulness). But the mind does not create – does not *invent* – the objects and the laws about them, about *it*, the same way it can think up unicorns. (Nor do the intersubjective dynamics between many minds bring those laws about as a social construction). It is only the “creative”, common interpretation of "mathematics being rooted in psychology" that turns AP8 against it. The popularity of this understanding is also due in part to Brouwer's position as the father of cognitive foundations for mathematics (Section 3.1). However, I suggest that it is not the only cognitive game in town. And that, in my suggested approach here, many of the issues that most mathematicians have with intuitionism and "created"

[16] To bring in Benacerraf and Putnam’s much later classification: “In general, the platonists will be those who consider mathematics to be the *discovery* of truths about structures that exist independently of the activity or thought of mathematicians. For others not so platonistically minded, mathematics is an activity in which the mathematician plays a more creative role” (Benacerraf, et al., 1984 p. 18).

mathematics need not arise. What *is* left is a different sense in which mathematics is *not* independent of us, of our psychology – because it pertains *to* it. In the case of MAC, this is because mathematics is concerned with a fundamental, idealized aspect of psychology – *any* psychology that may be, ours included. To conclude, AP8's quarrel is only with anti-realism, not with cognition per se.

Mathematics as discovered rather than created brings in a definite normative character. It is possible to get a law wrong, and then find out that, for example, counting to two twice does not result in five. This much may be true for BMAC just as it is for MAC. It does, however, go against the view of 'normative anti-psychologism' (Kusch), by which:

AP9) Psychology and mathematics are divided by the is-ought distinction.

As argued in Section 2.1, psychology as a science does have some access into the normative. But even if the view is accepted, AP9 poses no problem for the strengthened MAC, which takes out the worldly "is" component in the sense of having anything to do with what, for example, human psychology happens to be.

The sketch I have offered through this chapter illustrates the possibility of a cognitive foundation that conceivably could nevertheless withstand the familiar anti-psychologistic objections. This conception is what I shall term *cognitive realism*. It *is* cognitive, yet it is also all the following:

I) Fully spelled-out, explicit, no hidden assumption (ideally, when made ontologically rigorous).
II) Not vague or probabilistic but precise (as precise as we have come to know mathematics to be – and even more, if ontological rigor is to be successfully pursued).
III) Non-contingent (on our distinct makeup, on our evolutionary path, ideally not even on the laws of physics).
IV) Objective rather than personal (solipsistic) or merely intersubjective (a social construct). (As it is in particular not contingent on who we happen to be or what society we happen to belong to).

V) Not empirical, not (directly) about the physical world outside the cognitive system; and a priori, as in not known or justified through experience with that outside world.

VI) Abstract.

VII) Timeless.

VIII) Normative.

IX) Discovered rather than created.

In terms of the picturesque presentation of (Lakoff, et al., 2000) (Section 3.3) – it keeps to "the romance of mathematics". For cognitive foundations that do *not* keep in full to this sort of realism (e.g. varieties of BMAC), I suggest the term *neo-psychologism*.

2.7 Applicability

Mathematics extends beyond the ("intellectual exercise") sphere of pure mathematics, into the world. The need for a metaphysical foundation to ultimately account for that as well is an inherent challenge for a cognitive one. The simpler (though not at all simple) sense of mathematics having to do with the world, which I now address, is in its applicability to the world. We can add 2+2 apples (and predict the result), reason about a clock's or a crystal's symmetries (to reach real-world conclusions), etc. This applicability is representational (contrasting with the next section) in the sense that part of reality is represented, interpreted (modeled) abstractly, qua mathematics, and then answers to those mathematical laws that apply to the mathematical structure. Where does this applicability originate?

Logicism takes mathematics to *reduce* to logic (inheriting central aspects of its unique metaphysical nature). To the extent that logic constitutes "the (normative) laws of thought", this could be the source of the applicability of mathematics:

> "Why does arithmetic apply to reality?" the logicist provides the clear answer "Because it applies to everything that can be thought; it is the most general science possible." (Demopoulos, et al., 2005)

The applicability is not directly to the world but to the world as mediated through thought. But logic makes for a very particular form of thought, as a deductive activity upon pre-handed *propositional* knowledge. Computationalized cognition goes far beyond that, supporting *thought* in its complete generality – i.e. cognizing.

Mathematics, according to MAC, could be the true laws of thought writ large (or at least larger). The reason for its applicability to the world would then generalize accordingly; it applies to everything that can be cognized (and particularly it "is a priori true of the appearances" (Shabel, 2016)). (Note that for truly applying to everything that can be cognized, unconditionally, *cognition* should be interpreted in the most general computational sense, contra BMAC). Rather than reduce mathematics to logic, MAC thus takes mathematics to *extend* logic in a natural way (sharing central aspects of their nature).

Worldly applications of mathematics pose no inherent anti-psychologistic problem, then, if the mathematics can be interpreted as *about* representation. This includes most areas of inquiry that exhibit mathematical patterns (even if they are not completely mathematical, not lending themselves to a full model that captures everything of relevance). Scientific theories are not directly about the world itself the way it really is, independent of us in every way. Rather, they are always about the world as it is given to us, mediated through our cognitive framing – and necessarily so. It is just that we ordinarily see through that mediation, implicitly taking "the world" to be whatever we come to represent, however we come to represent it. A set of differential equations may *reflect* some population dynamics (for example) and thus provide an explanation of some fact of the matter (and generate predictions). But the equations need not *govern* those dynamics: They could not have been different in form (but only in empirical parameters), given the underlying fundamental physical laws. This is the same sort of applicability as in adding 2+2 apples, only employing more complicated mathematics. As long as the usage of mathematics is for *explanation*, there is a cognitive framing in the background, which it is in fact about. (All this extends to technology and engineering as well.)

A different sort of example is the immense mathematics that goes into the operation of a working brain and its construction (most of which we have probably yet to figure out as scientists or invent as mathematicians). Nature's great, 'blind' mathematician and engineer, Darwinian evolution, picks up on it, in a sense (Sloman, 2017). This discoverer isn't a cognitive agent; but the discovered mathematics *is about* cognitive systems, and so still in line with MAC.

The real problem would be with mathematics that *could not be* about cognition.

2.8 Mathematics in the Natural World

Our final anti-psychologistic objection concerns the appearance of mathematics – in nature:

> AP10) Mathematics is intrinsic to the infrastructure of the physical world and its operation – regardless of any cognizing beings.

Factually, AP10 certainly seems to be the case. It is an assumption that seems to stand at the core of the grand endeavor that is science. What room does this leave for cognition? All the previous objections were *challenges* for a cognitive foundation of mathematics that I suggest could conceivably be overcome. This last one, however, threatens to be an *insurmountable obstacle*; if AP10 is indeed true, then mathematics' *metaphysical* foundation could not be cognitive. Without getting into the necessitated deep metaphysics and the related literature[17], I only state, in brief and in the most general terms, why AP10 need not be conclusive.

What turns AP10 from a challenge to a threat is mathematics' applied involvement *in the non-cognitive* (or computational) – primarily with the fundamental laws of physics. Why would these be mathematical at all? The preliminary Kantian point (of the previous section) was that our theories are about the world as it is given to us, mediated through our cognitive framing. This point has suggested the source of applicability – of arithmetic or geometry, for example – to the world. In itself, this does not suffice for explaining the mathematicality of fundamental physics. But that cognitive mediation is extremely rich and non-straightforward. Further mathematical aspects of the world, thus, may be due not to how it "really is" (whatever that may mean) but to that cognitive framing of reality. To demonstrate and elaborate:

[17] For example, the discussions sparked by Eugene Wigner's widely-known "The Unreasonable Effectiveness of Mathematics in the Natural Sciences".

Consider the high point of physics becoming mathematical – Newton's law of gravity. The law is derived from (idealized) empirical observations. Even though it is mathematically expressed, conceptually, it may have been that it is completely "mathematically arbitrary". That this happens to be this particular equation, in this particular form, that gets it right physically, could be for no mathematical reason whatsoever. Yet its form, called an *inverse-square law*, repeats itself for many other fundamental physical laws (e.g. Coulomb's law of electrical forces). Could this have been just another, independent physical coincidence? That would be unlikely – and indeed it is not:

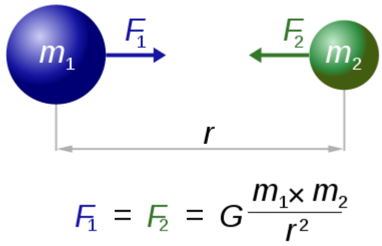

Fig. 1 Newton's law of gravity
(source: Wikipedia)

"The fundamental cause for this [form] can be understood as geometric dilution corresponding to point-source radiation into three-dimensional space" (Fig. 2). (Wiki-ISL, 2018)

Geometric framing thus matters, sneaking in, determining at least *something* about the physical law's mathematical form. And geometry itself, it is now agreed upon, is purely mathematical, a priori (and in our approach, possibly cognitive), rather than part of nature itself (though nature can happen to conform to some particular geometry – in a sense). Thus, mathematical reasons here play a role in a physical law being the way it is – even if it was discovered empirically. (This relates back to our earlier AP1 discussion on high-level facts being found through lower means). We can *stipulate* that the law could have

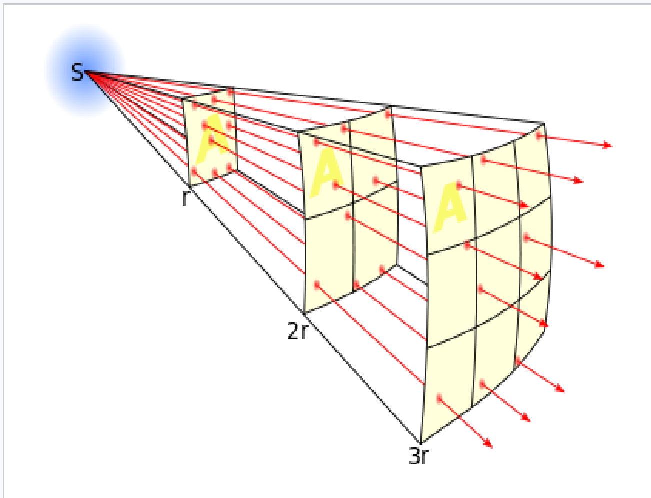

S represents the light source, while r represents the measured points. The lines represent the flux emanating from the source. The total number of flux lines depends on the strength of the source and is constant with increasing distance. A greater density of flux lines (lines per unit area) means a stronger field. The density of flux lines is inversely proportional to the square of the distance from the source because the surface area of a sphere increases with the square of the radius. Thus the strength of the field is inversely proportional to the square of the distance from the source.

Fig. 2 Geometric deduction of the inverse-square law been different in many ways, that it is physically determined but mathematically contingent – but that need not be the case; we simply do not see *why* that is not the case.

How far is the extent to which meta-physical facts are in fact deeper, mathematical necessities? This question cannot be answered in the a priori; hard scientific work is required. This tension between being part of the world and being an epistemic construct is nothing new to physics (e.g. in statistical mechanics). Theoretical reframing can be integral to advancements in the exploration of the laws of nature. And with it, as a general dynamic, things tend to become less mathematically arbitrary. That order arises from our framing – is always an option. And indeed, there are some current scientific works along such lines of "replacing physical ontology with epistemology" (though this is not the mainstream).[18]

The actual physical world may be meta-physically much richer than we perceive. (This is already suggested by science as a possibility, e.g. in quantum physics, dictating a deep break from our meta-physical, possibly innate intuitions). In such a case, the very starting point of the human endeavor to figure out the world's fundamental physical structure is a *cognitive construct*, already given to us through layers of cognitive pre-processing, reducing the meta-physics to what we do perceive and relate to. And to these layers, mathematics definitely applies (laws, invariants, etc.) – by my own cognitive conceptualization here, no need for a direct connection between mathematics and the world.

I thus suggest the possibility that, ultimately, it may *all be mathematical*; whenever a *law* is mathematical, the *reason* for it is actually mathematical (thus possibly stemming from our cognitive interface into the world). But that, of course, is only wild speculation. Deep metaphysical-meta-physical work is required here, beyond just the interface between mathematics and cognition (and hence beyond the scope of this dissertation). And without such an account, AP10 remains standing as a threat.

---

[18] To give one reference, Ariel Caticha (UAlbany) aims, "to explore the extent to which the laws of physics might reflect the rules for processing information about nature. More specifically the objective is to derive statistical mechanics, quantum mechanics, and general relativity as applications of entropic inference". (From his academic website.)

Even if the threat from fundamental physics can be dismantled, there may be other non-cognitive sources of mathematics in nature. A case might be made that evolution, too, discovers and makes use of mathematics in some sense (beyond the mathematics of information and control systems, which falls under MAC). But it is not clear how to make sense of that sense, and AP10 remains a threat in any case, so I leave it at that.

The very possibility of cognitive metaphysics along such lines therefore remains unsettled (though there still is room for a representational foundation and for providing insight through cognition, as I pursue in Part II). The existing cognitive metaphysical approaches that I review might fare differently. However, they are not always sufficiently explicit in addressing these classical, anti-psychologistic worries. And mostly, they all choose to bite some bullets that some philosophers of mathematics and many mathematicians feel uncomfortable with (spearheaded by non-realism).

3 Cognitive Approaches

Chapter 1 developed the general motivational narrative for a cognitive foundational approach. In doing so, it mapped, conceptually, some basic options for such approaches and how far they might go, metaphysically. Chapter 2 then questioned the familiar anti-psychologistic objections, and advanced the possibility of realism through MAC, another sense of grounding mathematics in cognition, as being *about* cognition. This does not yet provide for a systematic conceptual taxonomy that structures the space of different possibilities for a cognitive foundational approach. As anti-psychologism still rules, there are too few specimens to populate one anyhow.

To complicate things further: As previously discussed (Chapter 1.6), a cognitively informed approach need not *ground* mathematics in cognition. It may instead (P) ground mathematics in the physical alone (if not posit some separate reality altogether). Within the general foundational context and at this exploratory stage, the restriction here to a cognitive grounding may be methodologically *artificial*; of importance is the emphasis on the *interaction* between cognition and mathematics, more so than the philosophical stance as a preliminary dogma. Aaron Sloman, for example, takes mathematics to be in the world, a-cognitively. Evolution, as a "blind mathematician", then "discovers" mathematics and integrates it into biological mechanisms and in particular neural ones; much later, meta-cognitive mechanisms allow humans to discover mathematics in their own explicit terms. This conception does not put aside mathematicians' interaction with mathematics. But it does integrate its foundational exploration within a much richer context[19]. This physical-chemical-biological-evolutionary context is outside the scope of this dissertation (although it might be that such a scope is *necessary* for a foundation of mathematics). Thus, the work of Sloman and others along such lines is not covered here, though under a somewhat different framing, they are just as relevant.

Instead of a complete taxonomy, I suggest a pragmatically valuable, rough classification by the field the originator of a suggested approach is coming from. Underlying working philosophies, assumption, methodologies (etc.) determine the

---

[19] This context, for Sloman, is his Meta-Morphogenesis Project. See e.g. (Sloman, 2017).

nature of such a necessarily interdisciplinary approach to a considerable extent. This is true at least for the few existing approaches that I cover in this chapter. But there is one field that is privileged in a way, so let me bring it in in advance.

Mathematical cognition is the field within cognitive science that explores mathematical *thinking*. A priori, it does so with *no philosophical strings attached* beyond the ones customary for cognitive science at large. Some famous examples of its empirical discoveries are Piaget's developmental observations (e.g. that children at the "preoperational stage" develop "seriation," which is "essential for understanding number concepts… the ability to order objects according to increasing or decreasing length, weight, or volume" (Ojose, 2008)); the *SNARC effect* ("when presented with smaller numbers people respond faster with the left hand and when presented with larger numbers people respond faster with the right"); and *subitizing* ("rapid, accurate, and confident judgments of number performed for small numbers of items"[20]). Such findings (as always in cognitive science) come with related conceptualizations such as *the mental number line* (by which, as it would seem, number magnitude is represented spatially). A critical limitation of the field (with respect to mathematics at large) is its focus on the most concrete, directly applied, physically grounded end, which can easily be handled experimentally through familiar cognitive methodologies – chiefly numbers. This includes a general focus on what is common to all humans (and sometimes animals too). Only recently did works concerning more advanced mathematics, and particular to professional mathematicians, begin to appear (Pease, et al., 2013; Amalric, et al., 2016; Amalric, et al., 2018).

Interpreting empirical findings, saying what they mean or at least suggest, is an extremely tricky business. This difficulty may be true for cognitive science in general but has its own flavor and issues when it comes to mathematics. Where mathematics materializes implicitly in the world, it is particularly hard to tell what the subject is responding to (e.g. the number of apples or the continuous amount of apple surface). Thus, even the theory of "a number sense", one of the oldest and most robust theories

[20] Wikipedia. For a scholarly exposition see (Dehaene, 2011 p. 69), (Dehaene, 2011 p. 57) respectively.

to have emerged from the field, supported by many different findings, can still very much be disputed (Leibovich, et al., 2016; Núñez, 2017).

Relating empirical findings and their related conceptualizations to classical *philosophical* issues is *not* part of the science's purpose. These may certainly appear to bear on the philosophical issues, at least potentially. And sometimes (though not commonly), claims are made to suggest that some empirical finding attests not only to a more general principle of mathematical cognition but to deeper truths about mathematics itself, perhaps. Such suggestions, however, are just for stressing (or exaggerating) the potential importance, never the substance on which the works are evaluated and accepted.

I review central existing approaches to cognitive foundations, criticizing them and comparing them to my own viewpoint, in continuation of the previous chapters. I begin with one representative of mathematical cognition, a central figure in the field, who did attend explicitly to central philosophical issues (issues that may appear very different when approached from cognitive science, as in Chapter 1). I then continue to other main figures writing on the cognitive foundations of mathematics. But first, there is a brief historical prelude, an honorary mention of the historical cognitively inspired approaches that came from within mathematics itself.

3.1 Intuitionism & Constructivism

Within the philosophy of mathematics, a cognitive foundation/approach is almost synonymous with Brouwer's intuitionism (and, somewhat less so, with the various constructivist approaches that followed). Given its familiarity within the field, this section is brief and simply refers the interested reader to introductory expositions such as (Iemhoff, 2016; Bridges, et al., 2016; Posy, forthcoming). Brouwer's two acts of intuitionism, which form the basis of his philosophy and create the realm of intuitionistic mathematics, are as follows (included without elaboration):

1. Completely separating mathematics from mathematical language and hence from the phenomena of language described by theoretical logic, recognizing that intuitionistic mathematics is an essentially languageless activity of the mind having its origin in the perception of a move of time. This perception of a move of time may be described as the falling apart of a life moment into two distinct things, one

of which gives way to the other, but is retained by memory. If the twoity thus born is divested of all quality, it passes into the empty form of the common substratum of all twoities. And it is this common substratum, this empty form, which is the basic intuition of mathematics.

2. Admitting two ways of creating new mathematical entities: firstly in the shape of more or less freely proceeding infinite sequences of mathematical entities previously acquired …; secondly in the shape of mathematical species, i.e. properties supposable for mathematical entities previously acquired, satisfying the condition that if they hold for a certain mathematical entity, they also hold for all mathematical entities which have been defined to be "equal" to it …. (Brouwer, in (Iemhoff, 2016))

Intuitionism arose at a time when mathematics, growing into new, infinite realms, started to transcend the humanly graspable in an essential way, involving "less natural" idealizations, with regard to which opinions within the practice diverged. Mathematics was forced to reform and take a stance on new or generalized principles taken to underlie it (famously, the law of excluded middle or the axiom of choice) – or otherwise divide. This circumstance made the philosophy behind mathematics relevant to the practice itself (if not for the first time in the history of mathematics).

Brouwer's philosophy (with Kronecker and Poincaré as precursors) meant keeping mathematics – its objects as well as its principles – to something the human mind could relate to (a cognitive foundation that is *object-complete*; p. 18). In this basic sense, it was "cognitive". But the state of psychology at the time did not provide much *scientific* substance that could illuminate mathematics. Mostly, it was up to mathematicians' own ability to observe it introspectively. Brouwer, doing this *systematically* (from his particular theoretical-philosophical point of view), grounded mathematics in *phenomenology* (idealized, taken to be independent of human psychology[21]). Thus, "Brouwer's intuitionist school of mathematics and philosophy became more and more involved in what, at least to classical mathematicians, appeared to be quasi-mystical

[21] "[A]n idealized mind in which mathematics takes place already abstracts away from inessential aspects of human reasoning such as limitations of space and time and the possibility of faulty arguments." (Iemhoff, 2016)

speculation about the nature of constructive thought" (Bridges, et al., 2016). The scientific view of cognition today (as described in Chapter 1) is very different. Particularly, a phenomenological basis seems problematic, if not unwarranted, and in any case does not engage deeply[22] with the vast body of research that is rapidly accumulating, which is what this dissertation concerns.

It is not that there was no philosophical or phenomenological depth to that position (Posy, 2005). But that position entailed a deep *revision* of mathematics itself with respect to the main school of mathematics that was winning the day, becoming "classical" mathematics (see Section 1.1). The main, rich substance of intuitionism was of a technical-*mathematical* nature, revising mathematics in accord with the allowed principles. The main struggle was to reconstruct the body of mathematics in a way that is sufficiently applicable and rich (in terms of what it can prove) so as to be able to serve as a reasonable alternative. Then, presumably, the matter, within the practice, would come down to philosophical dispositions and mathematical intuitions. But dispositions and intuitions aside, the failures of Brouwer and his followers to achieve such a reconstruction is perhaps the main reason that "classical" mathematics, by and large, prevailed.[23]

The constructivist approaches that followed (Bridges, et al., 1987) took this tendency further. Disagreements about various specific principles entailed further splitting that brought about further rich technical-mathematical substance, without engaging with cognitive science with regard to the principles themselves.

This methodology reflects these alternative approaches as coming from *within* mathematics. There would certainly be value in examining their underlying assumptions from a modern cognitive perspective. However, from my own point of view, there is much to be done regarding the cognition that underlies mathematics *before* turning to revise mathematics itself if that turns out to be required. And so, the main technical body of these approaches is not particularly relevant to the rest of this

[22] Though certainly the connection *between* phenomenology and cognition is an important modern topic.
[23] Constructivism's restricted, more "down to earth" nature allows room for it as *part* of mathematics. It keeps finding its way into modern logic, computer science, and mathematics.

dissertation. Here, I just mention one central theme of intuitionism that *is* in the background:

Language

As quoted above, Brouwer strongly opposed the grounding of mathematics upon its language. In fact, Brouwer's view may seem like the reasonable default: In general, outside of mathematics, language is *about* objects and reality; it refers to them; reality stands independent of its linguistic reflection. It is the logic-grounded stance that proposed an alternative that is non-obvious, non-intuitive, and even contentious to begin with. Not without good reason, but still: Language supposedly playing such a different role in mathematics than it usually does, seems to call for and indeed *necessitate* a cognitive account of how it is that we come to use language in such a non-standard way. But instead, the logic-grounded approach, being anti-psychologistic, did the exact opposite, nullifying the tension.

3.2 Stanislas Dehaene

Dehaene (a world-class cognitive scientist) and the works he has been involved with constituted a substantial part of the mathematical cognition in the field's early stages, summed up in his classic (Dehaene, 1997). He is still a central figure, although the field has grown immensely since then. He has also co-edited a more advanced, up-to-date collection of works (Dehaene, et al., 2011). Although deep within the conventional empirical line of research, it explicitly suggests connections to the more philosophical related issues (especially Kantian themes). But it does not pursue them as such, and "processing" it philosophically would be a novel work that is outside the scope of this review. The former (Dehaene, 1997), however, does offer a chapter (9 – "What is a number?") explicitly on the philosophical outlook and implications stemming from mathematical cognition. Its general spirit was already reflected in Chapter 1 of this dissertation. I give a further selection of excerpts now (but I recommend reading the not-too-long stand-alone chapter itself).

The chapter begins with a factual focus on the brain – against a common conception of it as a logical deduction device. Neurons are the basic building blocks of its computation – *analog* computation. Interestingly, he relates this to empirical findings covered in previous chapters where comparison of numbers depends on their magnitude rather than their representation in the numeral system. This seems to suggest an analog coding of numbers (and "is not so easily implemented in a digital computer").

His suggestion, I suggest, is a good demonstration of the limitations of such cognitive research. There is nothing in such behavioral experiments to rule out a fundamentally symbolic representation of number, with extra computational layers that manage comparisons of magnitude, and that are perhaps what is accessed, influencing behavior. Modern neuroscience is now going beyond cognitive psychology, beginning to figure out the actual workings. Until this project is complete, however, all we can do is take empirical findings to suggest some possibility rather than another. [24]

How we happen to represent numbers need not carry any philosophical significance. However, the fact that we *do* – does. Criticizing the logical foundations from his cognitive angle, he takes *non-standard models* (similarly to my discussion in Section 1.1 and somewhat confused with Gödel's incompleteness theorem[25]) to show that:

> our "number sense" cannot be reduced to the formal definition provided by these axioms… The concept of number is primitive and undefinable. This conclusion seems implausible. We all have a clear idea of what we mean by an integer, so why should formalizing it be so difficult?...
>
> Ironically, any 5-year-old has an intimate understanding of those very numbers that the brightest logicians struggle to define. No need for a formal definition: We know intuitively what integers are. Among the infinite number of models that satisfy Peano's axioms, we can immediately distinguish genuine integers from other meaningless and artificial fantasies. Hence our brain does not rely on axioms.

This, of course, leaves open the question of how the brain *does* do all of that; how we know *intuitively* what integers are, how we can immediately distinguish them from "artificial fantasies". Could we use *that*, the underlying mechanisms, to produce a definition? Why not? He does not elaborate. The working assumption is that the brain

---

[24] He has more to say about that, but the point here is not to cover the debate, just give a taste of the issues at stake.

[25] The point here is not to belittle his own take. Quite the opposite, it is brought in as a positive example (against most texts in the field), trying to at least relate the topic of numbers to the deep question of how we could represent them.

does these things and that cognitive scientists would thus be the ones to figure out exactly how.

From there, he transitions to the question of the correct philosophy of mathematics. As already discussed, from his cognitive point of view, platonic intuitions should not be taken at face value. Scientifically they are considered "hard to defend" and "cannot be more than an illusion". Formalism is then ruled out as inadequate, failing to account for the focus on very specific systems. Finally, Dehaene arrives at Kant, Poincaré, and intuitionism, stating:

> Mathematical objects are fundamental, a priori categories of human thought that the mathematician refines and formalizes. The structure of our mind forces us, in particular, to parse the world into discrete objects; this is the origin of our intuitive notions of set and of number....
>
> Among the available theories on the nature of mathematics, intuitionism [e.g. Poincaré's, not particularly Brouwer's] seems to me to provide the best account of the relations between arithmetic and the human brain. The discoveries of the last few years in the psychology of arithmetic have brought new arguments to support the intuitionist view that neither Kant nor Poincaré could have known…
>
> Number appears as one of the fundamental dimensions according to which our nervous system parses the external world. Just as we cannot avoid seeing objects in color… and at definite locations in space… numerical quantities are imposed on us effortlessly through the specialized circuits of our inferior parietal lobe. The structure of our brain defines the categories according to which we apprehend the world through mathematics.

There may not be enough argumentative content here to turn the empirical findings discussed throughout the book into a philosophical support for such a position. But this is probably as an authoritative representative of the field as any – with regard to the philosophy that seems most compatible with the science too.

He then reaches the controversies within mathematics driven by such considerations:

> [T]his position should be clearly dissociated from an extreme form of intuitionism, the constructivism ardently defended by… Brouwer…

> It is certainly not for me to decide whether classical mathematics or Brouwer's constructivist mathematics provides the most coherent and productive pathways for research. The decision ultimately belongs to the mathematical community, and psychologists must confine themselves to the role of observer.

Thus, although he takes cognitive science to dramatically refute not only most philosophies of mathematics but mathematicians' own realistic intuiting of the objects they explore, he does not think this authority ranges over the actual practice too. The independent existence of the objects is an illusion of the mathematicians, and yet the finer aspects of mathematical reality as perceived are undistorted. No biases or illusions that cognitive science may shed light on are relevant, presumably; the practice will take after everything else. The separation between mathematics and its psychology is brought back in at this level, analogously to psychologists observing humans in general: they may demonstrate how a person misperceives a visual scene, but only when what the reality actually is – is agreed upon; psychologists are not an authority on *that*, on the true content.

As a matter of degree, I admit the possibility that mathematical practice just turns out to be strictly better than any other scientific method at cleaning after itself, through adjusting to contradictions, integrating with physical theories that are reflected in the world, etc. But I do not see a way to justify such a position in principle. A priori, cognitive science should have a say about the intuitions being formalized too (as discussed in Section 2.1), as is implied by what Dehaene says right after:

> Nevertheless, in my opinion both theories are compatible with the broader hypothesis that mathematics consists in the formalization and progressive refinement of our fundamental intuitions. As humans, we are born with multiple intuitions concerning numbers, sets, continuous quantities, iteration, logic, and the geometry of space. Mathematicians struggle to reformalize these intuitions and turn them into logically coherent systems of axioms, but there is no guarantee that this is at all possible. Indeed, the cerebral modules that underlie our intuitions have been independently shaped by evolution, which was more concerned with their efficiency in the real world than about their global coherence.

A fundamental motivation for mathematical realism, as discussed in Section 2.6, is the stability of its main body, so unique to mathematics, accumulating throughout history.

To support his non-realism, Dehaene thus pulls the spotlight away from that stability and then attributes whatever still stands to our cognitive core:

> Mathematics is not a rigid body of knowledge. Its objects, and even its modes of reasoning, have evolved over the course of many generations. The edifice of mathematics has been erected by trial and error. The highest scaffoldings are sometimes on the verge of collapsing, and reconstruction follows demolition in a never-ending cycle. The foundations of any mathematical construction are grounded on fundamental intuitions such as notions of set, number, space, time, or logic. These are almost never questioned, so deeply do they belong to the irreducible representations concocted by our brain. Mathematics can be characterized as the progressive formalization of these intuitions. Its purpose is to make them more coherent, mutually compatible, and better adapted to our experience of the external world… According to the evolutionary viewpoint that I defend, mathematics is a human construction and hence a necessarily imperfect and revisable endeavor.

This should be contrasted with the notion that cognitivism need not entail mathematical non-realism, as I have suggested in Chapter 2. Fundamental intuitions are formed, as the brain is shaped by evolution, within a unified reality. As part of that reality, I would suggest that there *must* be some globally coherent mathematical system. That system is what those intuitions are measured against; it is what they come to approximate. It is only when the intuitions as they *happen* to be are taken to be constitutive (BMAC), that there is no guarantee that this is at all possible. And even then, the practice's unified strive to find the most coherent global solution provides precision and refinement that might supposedly transcend those approximate innate intuitions. With respect to what would this precision be?

For Dehaene, the core in which the foundations of mathematics are grounded in is the "fundamental intuitions" that "as humans, we are born with" – that he takes his research to suggest are innate[26]. This philosophical position is therefore vested in the scientific findings and their interpretation and, ultimately, in the concrete cognitive fact of the matter. It may, instead, turn out that nothing (particularly) mathematical is innate, or at

---

[26] This is what he means by "are born with" here.

least that there is no number sense (which, as noted, is a real possibility debated in the current literature). This would debunk the proposed source of stability of that foundation.

I would instead suggest that whether evolution "prefers" and "manages" to implement fundamental intuitions innately or not is a pragmatic matter that should be of no consequence to such foundational debates (tracking the *source* of what, and how, mathematics is). This is not the *ultimate* source of the stability, of the intuitions' fundamentality. They might appear to us just as fundamental if they are learned through general mechanisms, picked up on from their all-encompassing integration into our environment and how everything behaves. What ultimately matters is the worldly evolutionary grounding, which assures that mathematics is not just an arbitrary invention. As he puts it:

> Throughout phylogenetic evolution, as well as during cerebral development in childhood, selection has acted to ensure that the brain constructs internal representations that are adapted to the external world. Arithmetic is such an adaptation. At our scale, the world is mostly made up of separable objects that combine into sets according to the familiar equation 1 + 1 = 2. This is why evolution has anchored this rule in our genes. Perhaps our arithmetic would have been radically different if, like cherubs, we had evolved in the heavens where one cloud plus another cloud was still one cloud!

And elsewhere:

> A quantitative representation, inherited from our evolutionary past, underlies our intuitive understanding of numbers. If we did not already possess some internal nonverbal representation of the quantity "eight," we would probably be unable to attribute a meaning to the digit 8… (p. 87)

Incidentally, suggesting a mental non/innate division within mathematics:

> I would like to suggest that [complex numbers etc.] are so difficult for us to accept and so defy intuition because they do not correspond to any preexisting category in our brain. (p. 87)

Thus, mathematical regularities are environmental regularities, stable though metaphysically non-necessary, perhaps. And when the mathematical is not a separate, abstract realm, but an adaptation to the environment, the famous “unreasonable effectiveness of mathematics in the natural sciences” might no longer seem like such a deep mystery:

> It may not be necessary, then, to postulate that the universe was designed to conform to mathematical laws. Isn’t it rather our mathematical laws, and the organizing principles of our brain before them, that were selected according to how closely they fit the structure of the universe? The miracle of the effectiveness of mathematics, dear to Eugene Wigner, could then be accounted for by selective evolution, just like the miracle of the adaptation of the eye to sight. If today’s mathematics is efficient, it is perhaps because yesterday’s inefficient mathematics has been ruthlessly eliminated and replaced. (Returning to his philosophical chapter, 9)

Ultimately, Dehaene squared the appearance of mathematics “in nature” with its supposed cognitive foundation as follows:

> The hypothesis of a partial adaptation of mathematical theories to the regularities of the physical world can perhaps provide some grounds for a reconciliation between Platonists and intuitionists. Platonism hits upon an undeniable element of truth when it stresses that physical reality is organized according to structures that predate the human mind. However, I would not say that this organization is mathematical in nature. Rather, it is the human brain that translates it into mathematics. The structure of a salt crystal is such that we cannot fail to perceive it as having six facets. Its structure undeniably existed way before humans began to roam the earth. Yet, only human brains seem able to attend selectively to the set of facets, perceive its numerosity as 6, and relate that number to others in a coherent theory of arithmetic. Numbers, like other mathematical objects, are mental constructions whose roots are to be found in the adaptation of the human brain to the regularities of the universe.

To summarize his suggested position in the terminology of the previous chapters: Mathematics is grounded in cognition; the mind does not merely “represent” numbers (etc.), which exist independently. The grounding, moreover, is metaphysically complete; there are no real non-mental objects (though the conception of mathematical ontology and logic is to be left for mathematicians to decide). The foundations for all the edifice of higher mathematical are the "fundamental intuitions", grounded in the mind, innately. As these have no absolute assurance, the approach is neo-psychologistic (p. 65). The implicit manifestation of mathematics in the physical world is not, strictly speaking, mathematical; mathematics is inherently manifested explicitly (subjectively and objectively), even if non-linguistically (as with animals). It is a version of the Kantian but species-relative, biological BMAC (p. 59). The grounding of mathematics is ultimately not in the cognitive alone but in its interaction with the physical (P-M), as assimilated through evolution.

3.3 George Lakoff & Rafael E. Núñez

George Lakoff and Rafael E. Núñez are two cognitive scientists who have written (Lakoff, et al., 2000). It is the most prominent book-length modern work (and to the best of my knowledge, the only one) that aims to establish a cognitive approach to the foundation of mathematics. Thus, it offers the most comprehensive, overarching coherent perspective (including a fairly proportionate sight on the topics that should be addressed, exemplifying many related themes) alongside the most detailed systematic, technical work for bringing that approach to life. Many find large parts of it to be misled or wrong in various ways, and accordingly, it has gained fame and influence mostly outside professional philosophy of mathematics circles. But at a more general level, it *can* serve as a "declaration of principles" of sorts and as a first representative demonstration of what it takes to account for mathematics in cognitive terms, for better such takes that hopefully follow. For this value and despite its flaws, I attend to it at some length in this section. I give a brief overview of the book, present its central themes, and characterize its philosophical position – all with respect to the context of this dissertation. I do not engage in a thorough critique of my own but make do with mentioning some existing critiques.

To characterize their approach from the outset first, according to the fields they come from would be *cognitive psychology*. (In fact, Lakoff generally does not work in mathematical cognition). This area is where high-level phenomena are explored, primarily through observation of behavior. Classically, this means a rather general separation from the neural level of the brain and the many imagining techniques that are taking much of cognitive science by storm (in the past decade or two). This work takes this trend further. Although (Lakoff, et al., 2000) places great importance on the empirical evidence on which it builds, there is little in it of what passes for direct empirical data in the cognitive science. The substance of their work mainly consists of observations and reflections upon mathematical concepts and their usage (e.g. linguistic), almost to the point of a phenomenological analysis. (My own work attributes more content to "the unconscious" than that). Centermost is how this is guided by their cognitive-science theoretical framing. They build on heavy-duty cognitive concepts, which are brought in from *outside* of mathematical cognition. This includes *metaphors*, where Lakoff had already done interesting (non-mathematical) work; and *the embodied mind* (Wilson, et al., 2017), in which they are both heavily invested. The reasonableness of their account is to be supported (as a starting point) by the reasonableness of the material they bring in from non-mathematical topics in cognition. Thus, this work as a cognitive one belongs to mathematical cognition in terms of what it tries to account for (the true mental mathematical content) but not (by itself) in terms of the concepts and methodologies on which it builds – which are general. This is unlike Dehaene, who bases his philosophical view directly upon the works in mathematical cognition.

3.3.1 Overview

The book consists of (roughly) six parts, which present the philosophy, methodology, some fundamental mathematical topics that are the most obvious candidates for any such cognitive foundation to account for, and one further, more elaborate case study. I now present these in brief:

Introduction: Why Cognitive Science Matters to Mathematics &
Part V: Implications for the Philosophy of Mathematics

This is the programmatic, purely philosophical part. The introduction is a manifesto that declares their position, which they set against "The Romance of Mathematics" (its standard, platonic, non-psychologistic conception). It may be taken as a relatively straightforward (or at least natural) view of mathematical ontology, its foundations, and the methodology by which it should be explored, including the empirical facts that should be taken into account – when viewed from the perspective of modern cognitive science.

Much deeper into the book, Part V revisits these philosophical themes, repeating them in a fuller manner, now relying on the technical content and demonstrations of their approach that the reader had already been through.

Part I: The Embodiment of Basic Arithmetic

The elementariness of arithmetic, its ties to actual life within physical reality, and its vast existing cognitive explorations make it the natural place to start, demonstrating upon it the technical contents of the authors' approach. They begin by reviewing basic cognitive experiments and facts and then interpret these for their own take. They introduced their "embodied mind" agenda and make it explicit by elaborately applying it to arithmetic and its laws (which it purports to account for, completely).

Part II: Algebra, Logic, and Sets

This part brings their neo-psychologistic approach to the most fundamental parts of modern, higher mathematics. It is a technical continuation of their approach (which is ultimately to be applicable to the whole of mathematics). However, dealing with extremely abstract realms stresses the tension between their grounded-in-life embodiment approach and the seemingly unearthly nature of mathematics, allowing them to give more substance to their non-standard take, to demonstrate its uniqueness, to confront it with the mainstream, known approaches.

Part III: The Embodiment of Infinity

From a cognitive point of view, infinity itself "must be biologically and cognitively plausible. That is, it must make use of normal cognitive and neural mechanisms" (p. 163). Beyond giving (what Lakoff and Núñez consider to be) such an account" of infinity, both potential and actual, its role in related areas is also examined: projective geometry, the real numbers and limits, transfinite ordinals and cardinals, and finally infinitesimals (including their own peculiar "granular numbers"). Various fundamental issues are dealt with along the way: the place of *processes* (rather than a static, timeless ontology) in the foundations of mathematics, *closure* as an ontology-generating force, and *quantifiers*.

Part IV: Banning Space and Motion: The Discretization Program that Shaped Modern Mathematics

The conceptualization of space is a cognitive topic at least as much as it is a mathematical one. Lakoff and Núñez posit a cognitive "Naturally Continuous Space" and then examined its “mathematization” (the former being able to accommodate various mutually inconsistent mathematizations). Various topics that arose are:

The role of infinity (as points are “infinitesimal”); the pre-mathematized intuition of points “touching” each other; numbers on the line (as numbers per se are conceptually distinct, an “application” of a separate mathematical structure); the quantified ε-δ as distinct from our true cognitive conceptualization of these matters (missing out on the conceptual meaning of approaching a limit); and the other accepted formalizations accepted by the community that may however be historically-contingent.

Part VI: $e^{\pi i} + 1 = 0$ – A Case Study of the Cognitive Structure of Classical Mathematics

Lakoff and Núñez embark on a long technical journey, aiming to give a conceptual account of perhaps the most famous equation in the history of mathematics. This means going through – in a sense *teaching* – trigonometry, e, i (and complex numbers in general; and their geometric interpretation), exponentiation of non-integer powers, logarithms… the whole required edifice.

Their purpose is to reveal the vast landscape of ideas behind a short equation that is “simply” about numbers; to expose its hidden meaning; and to "thereby demonstrate the real power of the approach… how the cognitive mechanisms described in the book can shed light on the structure of classical mathematics" (p. 10). However, there is very little here that would be new to anyone who has studied and understood all the above through standard (if well-taught) university courses. Its main value for us here, instead, is simply in posing the question of where, foundationally, everything else in mathematics (“practice”) aside from definitions and proofs belongs.

3.3.2 Central Themes

These are the central themes that reoccur throughout the book, underlying the technical details and the treatment of the many various topics mentioned above.

3.3.2.1 Full Cognitive Account

How much content is there to mathematical reality, reflected in the mind but not captured simply by the rules that objects and structures obey, as already represented in the logical foundations? One can pose – as a *cognitive* project, metaphysics aside – the task of giving a full account of all that is truly taking place in the mind. (Some level of normativity can be kept if, to exclude errors, the "correct" conceptualization is that which the mathematical community implicitly "agrees" upon, shares). This task is what Lakoff and Núñez call "mathematical idea analysis":

> Mathematical idea analysis… asks what theorems *mean* and *why* they are true *on the basis of what they mean*. (p. xv)
>
> What we are doing is making as precise as we can at this stage in cognitive science what was left vague in the *practice* of formal mathematics. A "formalization" of a subject matter in terms of set theory often hides the conceptual structure of that subject matter… What we are doing is making explicit what is implicit in the practice of formal mathematics: We are characterizing in precise cognitive terms the mathematical ideas in the cognitive unconscious that go unformalized and underscribed when a formalization of conscious mathematical ideas is done.… Formal Reduction… does *not* formulate otherwise vague ideas in a rigorous fashion… The original ideas are not kept [recall the discussion on *faithful representations*, Section 1.4]… the formalization is *not* an abstract generalization…The original ideas receive no mathematical idea analysis under the formalization. (p. 375)

According to them (as expressed particularly strongly in (Lakoff, et al., 2001)), mathematical idea analysis is a cognitive-science project par excellence. Mathematicians themselves are not obligated to pursue it – *nor do they even deal with*. Naturally, idea analysis *is* acute for mathematical *education*. Figuring out how mathematics is (correctly) conceptualized is undoubtedly central for the theory and the practice of teaching it. And so, they indeed touch upon mathematical education throughout the book. But the mathematicians' craft is completely different.

Methodologies aside, one may wonder how categorically differentiated what they speak of really is from what mathematicians and the mathematical community as a whole are doing. For example: "There is nothing in these axioms [ZFC] that explicitly requires sets to be containers… The axioms do not say explicitly how sets are to be conceptualized… Indeed, within formal mathematics, there are no concepts at all, and hence sets are not conceptualized as any thing in particular" (p. 145). What they consider implicit or even non-conceptual, however, can go a long way – sometimes even all the way – toward determining a notion – just not in terms they like or accept. When Mathematicians are defining stuff, modeling intuitive notions, putting them in terms of others – this is obviously *related* to their cognitive grasp if not underpinnings of the ontology. (Recall that Dehaene even takes that to be integral to what mathematics *is*). Moreover, the mathematicians themselves would consider what they are thereby doing as an *explicit* introduction (into mathematics) of the notions. At least this is an *ideal*. Mathematicians may concur that possibly more can and should be done; that further explication may be in place. In fact, they may cheerfully embrace (in the future) heuristic help from cognitive scientists, who might be able to go beyond what is consciously achievable for the working mathematician (as suggested in section 2.1). Not to see the large extent to which mathematical idea analysis is ultimately a *shared* topic and goal is, I assert, a false and extremely unproductive dichotomy.

I believe the ideal of *ontological rigor* that I have suggested (Section 2.3) is the right unification. The comparison brings up a general deficiency in the approach of (Lakoff, et al., 2000). The many "technical" cognitive accounts of mathematical topics that the book offers are schematic (which is to be expected given their very high-level "building blocks," which are the topic of the next sections). They may readily accept that this is just the beginning of the work that needs to be done. But this is not enough to put the project into a clear frame. When, exactly, has a sufficient cognitive account been given? When can we be sure that nothing further is going on subconsciously, unaccounted for, distorting the account? Without a principled answer, with no clear *standard of explanation*, it seems that we are not in a categorically better position than the mathematicians themselves when they feel that things have been made sufficiently-clear (a stability which may be shaken by future developments). The *major* merit their approach might miss this way is the possibility of rehabilitating upon a *cognitive* basis mathematics' unique character of absolute precision (and ultimately realism) – a central challenge to psychologism that *no such approach* can ignore. As noted, they just bite the bullet, rejecting "the romance of mathematics" – perhaps needlessly, as suggested and discussed in Chapter 2.

Deficiencies in their position aside, it does set an important direction for the philosophy of mathematics, going beyond "the idea that every subject-matter could be characterized in terms of an essence – a short list of axioms, taken as truths, from which all other truths about the subject matter could be deduced" (p. 109). What is needed is a fuller representational framework, in which axiom-style rules governing an ontology could be related (as they are!) to other conceptual contents concerning and even inherent to that ontology, and in which historical shifts in formalizations could be embedded and explored.

> If there are 'foundations' for mathematics, they are *conceptual foundations – mind-based* foundations. They would consist of thorough mathematical idea analysis that worked out in detail the conceptual structure of each mathematical domain, showing… just what the network of ideas across mathematical disciplines looks like.
> That would be a major intellectual undertaking. We consider this book an early step in that direction. (p. 376)

3.3.2.2 Conceptual Metaphor

Lakoff's work on *metaphors* is perhaps what he is best known for. A (conceptual) metaphor conceptualizes something in terms of something else:

> [it] is a cognitive mechanism for allowing us to reason about one kind of thing as if it were another… it is a cognitive mechanism that belongs to our realm of thought..: it is a grounded, inference-preserving cross-domain ["unidirectional"] mapping – a neural mechanism[27] that allows us to use the inferential structure of one conceptual domain (say, geometry) to reason about another (say, arithmetic). (p.6)

There is a great deal of collected linguistic data behind this, which is taken to show something deeper about the conceptualization:

> Affection, for example, is understood in terms of physical warmth, as in sentences like 'She *warmed* up to me', 'You've been *cold* to me all day', 'He gave me an *icy* stare', 'They haven't *broken the ice*'. As can be seen by this example, the metaphor is not a matter of words, but of conceptual structure. The words are all different (warm, cold,

[27] Of course, contemporary science does not yet offer a mapping of such a high-level cognitive mechanism to its actual neural underlying structure; this kind of talk is, well, highly metaphorical. In fact, it may well be that *no such neural mechanism* per se exist in *any* particular human being, as these are prone to errors, various influences from other faculties, etc. Idealizations may be sneaking in conceptual normativity (which Lakoff & Núñez strongly oppose).

> icy, ice), but the conceptual relationship is the same in all cases: Affection is conceptualized in terms of warmth and disaffection in terms of cold.
>
> This is hardly an isolated example:…
>
> Hundreds of such conceptual metaphors have been studied in detail. (p. 41)

This, incidentally, is a main source of the authors' widespread reference to the unconscious, which for metaphors is easily noticeable *once pointed too* – something which indeed some practitioners (of general language and likewise in mathematics) *may* be conscious of. But much more important philosophically, this is where *abstractness,* so essential to the nature of mathematics, is tied to cognition and abstract thinking at large:

> One of the principle results in cognitive science [so they claim] is that abstract concepts are typically understood, via metaphor, in terms of more concrete concepts (p. 39)…
>
> [Thus, metaphors'] primary function is to allow us to reason about relatively abstract domains using the inferential structure of relatively concrete domain. (p. 42)

That is, the concrete, which may be considered to *precede* the more abstract at the cognitive level (at least when it comes to stages of understanding), is given (within what they take for a "scientific metaphysics") a general *foundational precedence,* in fact a constitutive role, over the abstract. This is then taken to apply just as well to logic and mathematics:

Logical categories arise through a metaphorical mapping from bounded regions in space (where objects inside the region are the category members). From there, as conceptualized spatial containers tend to obey the laws of logic, these laws can be obtained: "Given two containers A and B and an object x, if A is in B and x is in A, then x is in B" ⇨ Modus Ponens, etc. Thus, "spatial logic is primary and the abstract logic of categories is secondarily derived from it via conceptual metaphor". (p. 45)

Lakoff & Núñez make do with finding a *central* concrete case that behaves as the abstract one does, central enough to be a contender for being "privileged" (usually from their embodied perspective, to which I turn soon). Classically, that contender would simply be conceived as a basic *instance* of the more abstract, more fundamental underlying structure. Much more than just a maybe-pointless debate regarding what comes first, the general metaphysical / epistemological status of *the abstract* is at stake here. This question may arise for some of their non-mathematical examples as well, for instance "Love is a [business] Partnership", with "profits from the business" mapping into "'profits' from the love relationship" etc. Such instances invite the suspicion that perhaps there is some deeper, shared abstract core. Put into non-philosophical terms: *why reason about the abstract at all*?

The classical answer, I think, would pursue the following general path. Abstract reasoning – particularly about *abstract ontology* – can be applied to many different realms that fall under its form. As such, rather than learn to reason about every one of them independently, one can generalize and become more knowledgeable of, and more efficient in mastering, *new* concrete realms – through the abstract. This view, however, requires *attributing a degree of reality to the abstract* at the theoretical level too (not just what the person as a subject thinks that exists).

To this, the authors zealously resist. Metaphors just take us from one realm to another, where truths in the origin transfer *per definition* into "truths" about the target domain: "It is the inference-preserving capacity of conceptual metaphor that permits these rules of inference to 'preserve truth'"(p. 137). Much, then, rests on this metaphorically grounded anti-realism: "Since [Cantor's proof] uses metaphor, it makes use of a cognitive process that exists not in the external, objective world but only in minds. Thus, the result that there are more real numbers than rational numbers is an inherently metaphorical result, not a result that transcends human minds". (p. 212, and similarly in many other places.)

The metaphysical and cognitive status of metaphors and the metaphor-system aside, their work does put a spotlight on, if not advance, the unified project of relating abstract thinking as a cognitive, empirical subject, to the details of mathematics. Moreover, their attention to basing each interaction of the mathematical contents with its applications – either outside of or within mathematics (e.g. numbers as lengths[28]) – motivates and advances a systematic account of the *applicability* of mathematics, integrating it into the very foundation of mathematics.

---

[28] "The Measuring Stick Metaphor" (p. 68)

This split regarding the outer-mathematical versus inner-mathematical source of a metaphor is, in their terms, the *grounding* metaphors versus *linking* metaphors. The linking metaphors take us deeper and deeper into abstract territories (according to them) – e.g. Boole's conceptualization of classes in terms of numbers and algebra. That work, however, would feel quite "familiar" to a reasonably aware mathematician. It is the call to account *foundationally* for all that is going on that is important.

The grounding metaphors, however, those grounding mathematics outside of it, in the real world, are more philosophically controversial. This is our next topic.

3.3.2.3 Embodiment

Embodied cognition is a popular, controversial buzzword. In general, "Cognition is embodied when it is deeply dependent upon features of the physical body of an agent, that is, when aspects of the agent's body beyond the brain play a significant causal or physically constitutive role in cognitive processing" (Wilson, et al., 2017). This concept relates to a broad collection of topics and research agendas. Some are philosophical conceptualizations, but some are concrete approaches that produce empirical contributions. Finding, for example, non-obvious bodily influences on conceptual processing ("when participants hold a pencil in their teeth engaging the muscles of a smile, they comprehend pleasant sentences faster than unpleasant ones" (Wiki-EC, 2018)); or sophisticated *interaction* between cognition and motorics, as with running toward a flying ball and catching it (which does not require cognition to track the full physics account but can make do with movement-related heuristics).

The authors' understanding of *the embodiment of the mind* is stated thus: "The detailed nature of our bodies, our brains, and our everyday functioning in the world structures human concepts and human reason. This includes mathematical concepts and mathematical reason" (p. 5). Whether truly the "profound insight from cognitive science" they take it to be or not, this is definitely not common knowledge, well established and agreed upon within the general cognitive community (and it is usually presented as an alternative to the mainstream). The gain, from their perspective, in basing cognitive *mathematics* on this as well is a unified, coherent mega-approach, a strengthened theoretical body, gaining importance and arguably a holistic support. The other side of this coin is that the controversiality of embodiment transfers to their account of mathematics (and alternatively, the failure of this approach for mathematics may reveal or sharpen the problematicness of their version of embodiment at large).

For Lakoff and Núñez, all of mathematics is to be ultimately grounded in the physical world and our activities in it. This grounding is achieved through the front line of applicability – the grounding metaphors. These are metaphors that are grounded in *non-mathematical, embodied* cognition – having to do with the bodily aspect of perception and action: "addition as adding [physical] objects to a collection, sets as containers" (p. 53), etc.

Basing the abstract on the concrete, the authors are led (as discussed) to choose *specific* concrete activities to ground mathematics in. Great significance is then placed on, first, applications that seem archetypal, and second, the dedicated underlying cognitive mechanisms, particularly those tailor-made for the concrete task. This approach is necessarily somewhat ad-hoc. At least the second, however, is bound to be a central issue for *any* cognitive approach. (Not just any *psychologistic* approach, as theirs, that views everything through the empirical findings, but any approach that simply refuses to renounce (the possibility of) the *relevance* of those mechanisms). For them:

> [T]he neural circuitry we have evolved for other purposes is an inherent part of mathematics, which suggests that embodied mathematics does not exist independently of other embodied concepts used in everyday life. Instead, mathematics makes use of our adaptive capacities – our ability to adapt other cognitive mechanisms for mathematical purposes. (p. 33)

The interaction between mathematics and our cognitive machinery, and the special status of mathematics, are extremely delicate issues, philosophically as well as scientifically. The authors' "concrete first" approach relies on a specific chosen carving of reality, of what the basic application is. Concrete, physical topics are easier to explore (e.g. measure), providing a better handle on for research. But I find it clear that the abstract may play an important, constitutive role in our cognition just the same. If we list the faults of a book (ordered by importance) and count them, do we really, necessarily, conceptualize it all through physical (material) objects first, with which our body interacts rather than just our mind, directly? An abstract structure shared between many tasks, both bodily and cognitive ones, may be what evolution adapts us for *directly*, through purely abstract reasoning, which is then applied to the concrete as needed.

3.3.2.4 Conceptual Blend

> A conceptual blend is the conceptual combination of two distinct cognitive structures with fixed correspondences between them. In mathematics, a simple case is the unit circle, in which a circle is superimposed on the Cartesian plane… When the fixed correspondences.. are given by a metaphor, we call it a *metaphorical blend*… an example.. is [1)] the Number-Line Blend, which uses the correspondences established by the metaphor Numbers Are Points on a Line. In the blend, new entities are created – namely, *number-points*, entities that are at once numbers and points on a line. (p. 48)

Some further fundamental examples are:

2) 2) "[T]he discretized number line, which is a conceptual blend of three conceptual domains – space, sets, and numbers." (p. 283)
3) Multiplication by negative numbers. According to them, it is conceptualized through *rotation*. This allows them to "ride the wave" of the familiar, well-researched, and robust cognitive topic of *mental rotation*; to assume an adaptation of the cognitive mechanism underlying it.
4) Last but not least are the natural numbers. These are based on our innate ability to subitize, extended through four different metaphors: object collection, object construction, segment measuring, and motion along a path. These different conceptualizations of numbers are combined to form a new, "metaphorical" type of object, which draws from each of its source domains. "Aside from the way they are mapped onto the natural numbers, these four source domains are not isomorphic: Object construction characterizes fractions but not zero or negative numbers, whereas motion along a path characterizes zero and negative numbers". (p. 80)

Grounding multiplication in rotation demonstrates the arbitrariness of this approach. When approaching multiplication from cognitive science, this connection may seem natural. But when approaching it from mathematics itself, the ad-hoc-ness of their approach is evident: for the multiplication by negative numbers (before stepping into the complex realm), there is no way to tell *rotation* apart from *reflection*, which are nonetheless conceptually distinct. Their account over-determines what is taking place (quite possibly at the cognitive level, too), adding extra content that is mathematically irrelevant (akin to sets that are taken to represent numbers). And it is doubtful that rotation is what, say, children come to have in mind at first. The space through which the rotation supposedly takes place is not part of what integer numbers *are*.

Their particular choice of foundational building blocks and how they use them to account for the mathematical objects aside, the general, underlying point is one of their most important ones. Ontology should be accounted for not just by the laws the objects obey but by the processes that underlie this, and that determine how the objects are *used*, applied both in and out of mathematics. Formalizations can hide such details, giving a more *useful* characterization (for mathematical *practice*) – at the price of hiding some of the substance of the full, *actual* foundations.

Metaphorical blends are a central concrete manifestation of the general "Full Cognitive Account" doctrine. "Many of the most important ideas in mathematics are metaphorical conceptual blends. As will become clear in the case-study chapters, understanding mathematics requires the mastering of extensive networks of metaphorical blends" (p. 48). The actual content that unfolds throughout the book may seem quite familiar – if perhaps in retrospect – for those who know and have mastered the mathematical concepts. But the fine details of inner-mathematics applicability, how mathematical structures are superimposed, how things are meshed together – things that are all part of what mathematics itself is for us – seems to be missing from the classical (e.g. set-theoretical) foundational accounts.

3.3.3 Philosophical Stance

As the authors' philosophy of mathematics is very different from the mainstream approaches, they place them – platonism, logicism, formalism, and set-theoretic foundations – into that one general opponent of "the romance of mathematics". (They hardly ever mentioning intuitionism, and never nominalism). In contrast that "mythology", their own position can be unfolded as follows:

Platonism – as a philosophy – they take to be *non-scientific*; a matter of belief in a mathematics that "transcends" *human* mathematics (perhaps on par with theological beliefs and the likes). Beyond their general position, they presumed to give the following specific argument:

> [E]ven on the basis of mathematics itself – without any scientific evidence – the claim that transcendent mathematics exist appears to be untenable. One important reason is that mathematical entities such as numbers are characterized in mathematics in ontologically-inconsistent ways [points on a line; set-theoretically; through combinatorial game theory (by Conway and others)]… Since transcendent mathematics takes each branch of mathematics to be literally and objectively true, it inherently claims that it is literally true of the number line that numbers are points [etc.]… None of these branches of mathematics has a branch-neutral account of numbers…. This is an inherent problem in the philosophical paradigm of transcendent mathematics… But human mathematics is richer [!!]… It simply has many ontologically distinct and incompatible notions of number… To make sense of the ontology of number in mathematics, one has to give up on the idea that mathematics has an objective existence outside human beings and that human beings can have correct knowledge (via proof) of transcendent mathematics. (p. 342).

This "ontological inconsistency", as I see it, is mainly with respect to their straw-man. They would obviously have to work harder than they do to tell the world something new about inconsistencies in mathematics – even if they are not logical, deducible inconsistencies. More generally, the issue might instead be unrelated to transcendence. If cognition can square things, why can't the platonic realm offer a basis that can be tied to different realms in different ways? Rather than having “many ontologically distinct and incompatible notions of number”, numbers may be (as they are standardly taken to be) a sui generis structure that can be *instantiated* by various systems, mathematical ones included. That would simply be what human cognition perhaps does manage to *represent*, then (as in Chapter 1); the general *notion* of a number. There are valid scientific reasons not to take for granted the platonic realm, but its poverty with respect to the mental is not one of them. Without further argument, any added richness to the mental that *in principle* could not be part of the mathematical realm as a platonic one could only be due to inconsistency (or more generally, incoherence) alone. Weeding it out is in fact an essential component of mathematics and its progression.

This straw-man aside, the cognitive approach (as discussed generally in Section 1.4–1.6) does revive, unabashedly, the grand goal of a *metaphysical* foundation. What mathematicians have in mind (e.g. in “2” or in “ℕ”) cannot be different from what they *mean* (by “2” or by “ℕ”). This is in contrast with other representations ({{},{{}}} or first-order Peano arithmetic) that might fail to represent the intended. The main question is whether the mental is the representational gateway, or the thing itself. Lakoff and Núñez's position, as empirical scientists, is that of mathematics as *human* mathematics. All is empirical, and hence so is their *realism – physical*, and cognition driven, through metaphor:

> [I]n each of the 4Gs [the metaphors that ground arithmetic], numbers are things that exist in the world... These four metaphors thus induce a more general metaphor, that *Numbers Are Things in the World* [my emphasis]… [this] has deep consequences… [for] the widespread view of mathematical Platonism. If objects are real entities out there in the universe, then understanding Numbers metaphorically as Things in the World leads to the metaphorical conclusion that numbers have an objective existence as real entities out there as part of the universe. This is a metaphorical inference from one of our most basic unconscious metaphors. As such, it seems natural. We barely notice it. Given this metaphorical inference, other equally metaphorical inferences follow, shaping the intuitive core of the philosophy of mathematical Platonism: [truth-value realism, and truths as discovered rather than created; objects as not products of the mind (and many more on page 350)]… What is particularly ironic about this is that *it follows from the empirical study of numbers as a product of the mind that it is natural for people to believe that numbers are* not *a product of mind!*" (p. 80)

In the terminology of Chapter 1, they establish a foundation that *could* perhaps have been interpreted as a *representational*, *meta*-mathematical (but non-mathematical) foundation. However, naturalism pushes to accept it instead as the true, only real ontology – the foundation as a *cognitive* (and object-complete) foundation. A more general metaphor (somehow "induced") that they further suggest explains why it seems to mathematicians that they are merely representing. Getting into the cognitive details of what makes mathematicians attribute objecthood, and examining parallels between perceived mathematical objects and physical ones (which are also mediated through cognitive processes), is critical for any cognitive approach to mathematics, even if their particular account, which "follows from the empirical study", can be debated.

Their position, then, would best be described as *neo-psychologistic*. Naturally, they are psychologistic even regarding logic itself: "The four inferential laws we will discuss are Container schema versions of classical logic laws: excluded middle, modus ponens, hypothetical syllogism, and modus tollens. These laws, as we shall see, are cognitive entities and, as such, are embodied in the neural structures that characterize Container schemas. They are part of the body. Since the do not transcend the body, they are not laws of any transcendent reason" (p. 135). Rather than logic depending just on our cognitive makeup, for these laws as cognitive *entities* it thus matters what one has happened to study: "Of course, Boole's metaphor does not exist for most people. In most people's ordinary nonmathematical understanding of classes, there is no empty class and no such 'truth'" (p. 367). And if for logic, then even more so for mathematics: "$\aleph_0 + \aleph_0 = \aleph_0$ is true only by virtue of accepting those metaphors" (p. 368).

As I have argued in Section 2.6, a cognitive approach to mathematics need not commit to ontology being in any sense *created* by the mind. A cognitive realist might suggest instead that "$\aleph_0 + \aleph_0 = \aleph_0$ is true by virtue of those metaphors themselves" – whether or not someone – anyone – accepted them.

At least for the authors too, "the potential for mathematics... is a human universal" (p. 351). But they do not take this further into a holistic conception of mathematics. There is no *meta-mathematical corral* (p. 10), not even a cognitive one, capturing in any sense all the mathematics that could ever be – as a unified whole; all possible concepts and what their "acceptance" *would* entail. Instead, mathematics is "open-ended. Can be extended to create new forms" (p. 379). And there certainly could not be such a *mathematical* foundation. As they differentiate metaphysically between mathematics and its deeper cognitive underpinning, "*Mathematics cannot adequately characterize itself!*" (p. 348). There is no novel cognitive foundational analog to the *mathematical* logical/set-theoretic foundations they criticize.

Under their physicalist conception of mathematics, which is contingent upon our cognitive makeup, there is some room is for *normativity*. Mathematics is ultimately grounded in our bodies and reality, with metaphors simply "preserving truth" by fiat (as these map into realms to which they do not attribute independent reality anyhow). Thus, normativity comes from nature, from the sources of the metaphors (e.g. abstract containers obey some laws of physical ones). Much more generally, this is where *all* of mathematics' unique properties come from:

> Mathematics is a mental creation that evolved to study objects in the world. Given that objects in the world have these properties, it is no surprise that mathematical entities should inherit them. Thus, mathematics, too, is *universal, precise, consistent* within each subject matter, *stable* over time, *generalizable*, and *discoverable*. The view that mathematics is a product of embodied cognition – mind as it arises through interaction with the world – explains why mathematics has these properties. (p. 350)

For example:

> The stability [within the mathematical community] of embodied mathematics is a consequence of the fact that normal human beings all share the same relevant aspects of brain and body structure and the same relevant relations to their environment that enter into mathematics… (p. 352)
>
> Once we learn a basic cognitive mechanism, it is stable in each of us". (p. 355)

With this grounding, they stop short of *social* constructivism. Mathematics is not arbitrarily shaped by history and culture alone – because our mental and physical makeup matters too.

Constituted upon a cognitive foundation, mathematics cannot be literally in the fabric of the world. And so, like Dehaene, they naturally opt for the “Kantian” direction of understanding physics as perceived through our human conceptualizations (numbers are not in the actual path of the ball, etc.). Mathematics is useful, *applicable* to the world – because cognition is, and particularly because the mathematical metaphors are grounded in application. The *effectiveness* of mathematics is a "tribute to evolution and culture… That effectiveness results from a combination of mathematical knowledge and connectedness to the world". (p. 378)

I, again, would stress that this explanation may account for the *specific applications* in which the metaphors are supposedly grounded to begin with, while the reach of mathematics’ applicability grows so much further, to *novel* applications in latter scientific endeavors – through the abstract.

3.3.4 Reception

Lakoff and Núñez do not come from the philosophy of mathematics (as a profession), and their book is not quite directed toward those circles. Furthermore, it suffers from many stylistic, technical, and scholarly problems, which gravely hinder its influence upon the field, where it has thus received relatively little attention.[29] Meanwhile, it has received much attention outside the field (including among some mathematicians too). This demonstrates, I believe, the wider interest among mathematicians and cognitive scientists alike in the philosophy and foundations of mathematics that is not satisfied by the offerings of the current logical foundations and particularly by its disconnection from the mental.

The book's flaws notwithstanding, it is indeed "an unprecedented attempt to base advanced mathematics on the workings of the brain"[30]. Hence, I have brought it here at length, for the sake of the context of this dissertation. However, I have omitted providing a thorough critique for its own sake, and instead now briefly include some central reviews that have appeared (and with which I generally agree).

(Davis, 2005) is an excellent and very critical review, particularly technically. In common with much of what I have focused on, it criticizes the book on other various points as well, including:

- Metaphysical sloppiness regarding how orderly the physical domain is to begin with, if numbers are to be constituted upon it; and other examples where it seems that their account is not truly naturalistic but implicitly still platonist.
- A general lack of cognitive evidence.
- Confusions between doing cognitive science and doing prescriptive mathematics.
- Their "woefully inadequate" account of infinity.

(Goldin, 2001) is another very critical review, extremely dismissive, that notably appeared in *Science*, and accordingly, it criticizes how very little of modern cognitive science actually enters the authors' account.

(Madden, 2001) is important for being written by a mathematician – for mathematicians (appearing in the AMS). It locates the project nicely within this vague sphere of other related works. It raises the particular issue of the many different conceptual ways from which different mathematicians might approach the very same mathematical structure

[29] (Wagner, 2017) is one example of an exception, that does take on a deeper (if also critical) engagement with the book.
[30] Reuben Hersh, book cover.

(something which does not receive the attention it deserves in the book – as the authors postulate their particular accounts of the various mathematical ontologies as *the* true ones).

(Voorhees, 2004) is a later review that thus had the benefit of considering the other various reviews that had already appeared, including Lakoff and Núñez's own response (2001) to a short review by Bonnie Gold.

Finally, one notable exception to the generally critical line of reviews, is (Schiralli, et al., 2003). Appearing in a mathematical *education* journal, this, accordingly, is its focus. It claims that the book "has had a significant impact on the mathematics education community". It aims to sketch a plan for salvaging the general direction of the book from its many flaws. However, it brings back a categorical distinction between mathematics and mathematical thinking and seems much more relevant to education than to philosophy.

3.4 Giuseppe Longo

Giuseppe Longo is another central modern figure who has written extensively on the cognitive foundations of mathematics. Coming from theoretical computer science, he trades the cognitive expertise and prestige of Dehaene, and Lakoff and Núñez, for a much better-informed position regarding the history, the philosophy, and the logical foundations of mathematics, as well as much better acquaintance with modern, higher mathematics itself[31]. Various parts of his position are expounded in a long list of papers (Longo, 1999; 2002; 2005; 2005; 2007; 2010), (Longo, et al., 2010), and others, and more unitedly in a book (Bailly, et al., 2011)[32] (though of a scope much wider than ours here, also ranging over physics and some biology). I include central points of his position through a selection of excerpts (from the book, unless specified otherwise). Some of his themes are common to cognitive approaches in general, as has been presented in Chapter 1, shared with the previous figures I have spoken of.

---

[31] He has co-written the book "Categories, Types and Structures. Category Theory for the working computer scientist"; served as the editor of the journal "Mathematical Structures in Computer Science"; etc.

[32] The book is co-written with a physicist, but the content on mathematics and cognition, which I address, is mainly Longo's.

Opposing platonism, he calls to set as primary the connection between mathematics, the physical world, and us and our cognition. The approach is holistic-scientific.

> [M]athematics is one of the pillars of our forms of knowledge, it helps to constitute the objects and the objectivity as such of knowledge (exact knowledge), because it is the locus where "thought stabilizes itself"; by this device, its foundation "blends" itself to other forms knowledge and to their foundations. Moreover, the conceptual stability of mathematics, its relative simplicity (it can be profound all the while basing itself upon stable and elementary, sometimes quite simple, principles) can provide the connection which we are looking for with the elementary cognitive processes…
> The problem of the cognitive foundations of mathematics must therefore be analyzed as an essential component of the analysis of human cognition. (Bailly, et al., 2011 p. 20)

Relying on a very broad notion of foundations, the project as he sees it is to

> produce the constitutive history and the evolving, cognitive and historical (including cultural), foundations of mathematics". (Longo, 2010)

3.4.1 Incompleteness

Longo draws considerable justification for the move towards cognitivism from the limitations of the logical foundations approach (a bit from non-standard models, but mostly from the incompleteness results), whereas

> We humans are absolutely not constrained to reason with no reference to meaning… our rigor is not simply formal/linguistic. (Longo, 1999)

He critiques the logical foundations along similar lines to chapter 1, and (together with his co-author)

> [introduce] a distinction between "principles of proof" and "principles of construction" which we develop in parallel between the foundations of mathematics and of physics as well as within these disciplines. We will then understand the great theorems of incompleteness of formal systems as a discrepancy (a "gap") between proofs and constructions… [that] prohibit[s] the reduction (theoretical and epistemic) [of the latter to the former] (or also of semantics – proliferating and generative – to strictly formalizing syntax). (Bailly, et al., 2011 p. vii & 2)

This includes an analogy of mathematics and its logical foundations with physics and *positivism* as a restrictive philosophy, akin to chapter 1's psychology as restricted to behaviorism (which was indeed influenced by positivism).

The two levels (of proofs vs. constructions) are not completely disjoint:

> [I]t is the coupling and circulation between these two levels that make this articulation between innovative imagination and rigor which characterizes the generativity of mathematics and the stability of its concepts. (p. 4)

Thus, contra logicism:

> [M]athematics has a logical as well as a formal foundation (a distinction will need to be made here), but it is in fact a "three-dimensional" construction. It constitute itself within the interactions of the logical and totally essential "if . . . then" (first dimension), of perfectly formal, even mechanic calculus (second dimension), but also in a third conceptual dimension, these constructions of (and in) time and space, which mingle it, even more so than the two others, with the different forms of knowledge. (p. 22)

All in all, against Frege's attack on psychologism:

> Frege proposes a philosophy centered upon a very inflexible dogma… according to which mathematics has no psychologico-historical or empirical genesis. (p. 21)… the insufficiency of a sole logico-formal language as the foundation of mathematics… brings back to the center of our forms of knowledge a constitutive mathematics of time and space (p. 33)

3.4.2 Signification

That third dimension of mathematics, bestowing *meaning* upon the logico-formal, is the *signification*, through *gestures*:

> [M]athematical gesture [is]… in short: a mental/bodily image of/for action… (p. 58) [S]ignification is constituted by the interference of signal with an intentional gesture… gesture, which begins in motor action, set the roots of signification between the world and us… The neuron … reacts with an action, a gesture at its scale, with its internal and external mobility; at its level, this reaction is meaning… (p. 61)

> [T]he neural network also changes and the net of networks, and the brain as well, in a changing body. This is the modified activity of this entanglement and the awfully complex coupling of organization levels which makes significant the friction between the living and the external world… perception itself is equivalent to the difference between an active forecast and a signal… For this reason, perception leads to signification… there is no meaning without an ongoing action. (p. 62)

A symbol is "a synthetic expression of signification links":

> Each symbol, each sentence has a huge "correlation length" in the space of present and history… This length describes the possible distance of causal links: each word is almost physically correlated, by an individual or a whole community, to a huge set of words, acts, gestures, and real-life experiences… There is no metaphysics of the ineffable, but rather a concrete, material and symbolic reality of phylogenetic, ontogenetic, and cultural complexity of the human being and his language… This takes us far away from "thought as a formal computation." Yet, mathematics is symbolic, abstract, and rigorous. (p. 62-3)

3.4.3 Between Cognition and the World

With meaning to be found in the friction between the living and the external world, the applicability of mathematics to the physical world thus goes through a (non-Kantian) form of the biological BMAC (p. 59). In brief:

> [Mathematics] is conceived on the "interface" between us and the world, or, to put it in Husserlian terminology, it is "designed" on that very "phenomenal veil" by which, simultaneously, the world presents itself to us and we give sense to it." (Longo, 2002)

Being central to (Bailly, et al., 2011)'s topic, this theme is developed at quite some length there (much more than I will offer in this dissertation).

Mathematics is thus *imposed* upon the world:

> "Properties," as we render them through intersubjectivity by words, are not in themselves isomorphic to absolute facts that are "already there,"… by our active gaze, in our exchange with others, we propose a structure with hints of a reality which is there, as unorganized frictional matter. Thus, through language, pictures, gesture, we

> unify certain phenomena, we draw contours upon a phenomenal veil, which is an interface between the world and us. (Bailly, et al., 2011 p. 23)

In general:

> Mathematics is the result of an open-ended "game" between humans and the world in space and time; that is, it results from the intersubjective construction of knowledge made in language and logic, along a friction over the world, which canalizes our praxes as well as our endeavor towards knowledge. It is effective and objective exactly because it is constituted by human action in the world, while by its own actions transforming that same world. (Longo, 2010)

As for objectivity,

> the objectivity of the constructed, of the concept, of the object, lies in the constitutive process, which is itself objective. (Bailly, et al., 2011 p. 24)

The constitutive role of cognition stresses the analogy with *physical* objects (as conceptualizations). Putting these on par (following Gödel):

> [brings] the question of a mathematical ontology closer to that of an ontology of physics… However, the difference… is given by the understanding of the object as constituted; it is not the existence of physical objects or of mathematical concepts that is at stake, but their constitution, as *their objectivity is entirely in their constitutive path*. It is thus necessary to take Gödel's philosophy… and to turn it head over heels, to bring it back to earth: one must not start "from above," from objects, as being already constituted (existing), but from the constitutive process of these objects and concepts. This requires a non-naive analysis of the object and of physical objectivity, as well as a non-passive theory of perception. (p. 28)

(Taking the cognitive *constitution*, of what we come to *take* for objects (physical and mathematical on par), as the foundational *starting* point, will be central to my own approach, in the chapters to follow).

In the terminology of chapter 1, mathematics, explicit, is cognitive, but its meaning is grounded in the interaction with the physical – (P-M):

> Mathematics is thus not a logico-formal deduction, nicely stratified from these axioms of set theory that are as absolute as Newton's universe, but is structurations of the world, abstract and symbolic, doubtless, yet not formal, because significant; its meaning is constructed in a permanent resonance to the very world it helps us understand. They then propose collections of "objects" as conceptual invariants, of which the important thing is the individuation of the transformations which preserve them, exactly like (iso-)morphisms and functors preserve categorical structures (properties of objects of a category). (p. 35)

Permanent resonance to the world, the *on-going interaction* with the physical, precludes this version of BMAC from being a Kantian, a priori one. This is unlike Dehaene's innateness-centered version, where the interaction is on the evolutionary scale, its result imposed upon the individual, and hence in this sense "a priori".

3.4.4 Proofs

Invariance is key, not only for the constitution of objects, but for proofs. Within the general body of knowledge, invariance (and stability) determine the special status of mathematics, which proofs reflect:

> [M]athematics is *symbolic*, *abstract*, and *rigorous*, as are many forms of knowledge and human exchange. But it is something unique in human communication, based on these three properties, since it is the *place of maximal conceptual stability and invariance*. … Invariance is imposed on proofs. This can even constitute a definition of mathematics: as soon as an expression is maximally stable and conceptually invariant, it is mathematics. But let us be careful, we use 'maximal' rather than 'maximum' because we aim to avoid any absolute (p. 64)

Thus, the stability of proofs and the *accumulative* nature of mathematics is achieved by *construction*:

> [P]roof-theoretic invariance is more a practical invariant than a just a formal invariant; i.e. it is constructed in the conceptual, scientific praxis of mathematics, as part of human reasoning. (Longo, 2002)

That mathematics is maximally stable is certainly true – *extensionally*. However, I would question whether this is all that there is to the essence of mathematics. The nature of his related distinction between "absolute" stability and his "maximal" stability (a distinction which is crucial for him and appears in many places), is not clear. It seems to suggest, for example, a rather smooth transition between what is mathematical and what is not (as opposed to differentiated categories). This may include perhaps even fluctuations in time. With no *absolutes*, can "maximally stable" be a permanent status? Or might something that was once "maximal" and hence mathematics, later on, with the arrival of *more*-stable things, cease to be so? Does it often happen in history that the community "changes its mind" regarding whether something is mathematics or not? Perhaps some such interpretation of the history of mathematics *is* possible. But much more should be said, given that it seems (as it is often taken to be) that mathematics is unique in some deeper, metaphysical sense. This, as discussed in chapter 2, is an important part of why non-psychologism has caught on, philosophical arguments aside.

3.4.5 Intersubjective

Stability arises from the practice as a whole, which is inherently *intersubjective*. This is part of what signification includes. This is part of how the true historical-cultural account of the development of mathematics is taken to be fundamental to the foundational project:

> [T]hese [technical] developments [of intuitionism, following Brouwer but not Poincaré] have undergone, on the one hand, a complete loss of any sense because of the formalization of intuitionist logic by Heyting and his successors… and, on the other hand, by the philosophical impasse of the Brouwerian solipsism and language-less mathematics…: such ideas are completely opposed to the constitutive analysis we propose here, which refers, in a fundamental way, to the stabilization of concepts occurring within shared praxis and language in a human communication context. (p. 80)

Taking as constitutive the stabilization within the shared practice seems to entail[33] a *non-revisionary* approach towards mathematics, akin to the Shapiro-Hintikka "mathematicians basically know what they're doing" approach (chapter 1). Again, I would suggest that modern cognitive science does not allow for this. Human biases

[33] I don't recall an explicit allusion of Longo's to the issue of revision of mathematics. It is certainly not central for him, regardless of any other philosophical disagreements he may have e.g. with Brouwer on the principles by which to actually revise.

(constituting "human reasoning"), for example, can be quite stable. Cognitive science is about going deeper than that. Its potential conflict with what mathematical practice comes to produce cognitively-"naïvely" through its own means cannot be ignored, in principle (even if the time is not yet ready for drawing *ontological* conclusions). One might be drawn into the conflict when beginning to account for higher mathematics systematically in cognitive detail, but Longo does not usually get that far.

The history of mathematical concepts and their evolution may certainly help shed light on our cognitive makeup and how it relates to them. But is what the practice produces by stabilizing on actually *constitutive*, then? I, instead, would suggest (more realistically, and not very originally) that to the extent that an aspect of mathematics is culture-dependent, contingent in this sense, it is not *fundamentally* mathematical; Whatever is fundamental *drives* (among other things that also influence) the dynamics of the practice to stabilize upon certain concepts etc. *If* language (logic included) is inherently intersubjective (which is reasonable but perhaps not necessary), and *if* it is a fundamental dimension of mathematics (unlike for Brouwer, for example), then the *fundamental intersubjectivity* of mathematics indeed follows.[34] But fundamental intersubjecticity is stronger than just being shared, for example due to corresponding cognitive mechanisms shared among the species (and maybe with other species too). And the constitution of stability, for Longo too, finds support in actual human cognition:

> The capacity to forget in particular, which is central to human (and animal) memory, helps us erase the "useless" details; useless with regards to intentionality, to a conscious or unconscious aim. The capacity to forget thus contributes in that way to the constitution of that which is stable, of that which matters to our goals (p. 42)

This is part of his grounding in the particular biological mechanisms (BMAC). Such details may indeed influence the path the practice comes to take. But I find no support here (and in similar such discussions) against grounding mathematics in the non-biological, abstract cognition MAC, which could keep to mathematical realism as

[34] This depends on the particular sense of "language". A general-enough one, e.g. a Fodorean "language of thought", may perhaps allow non-"linguistic" animals to perceive mathematical facts etc.

classically conceived. Forgetting is not just a human (or animalistic) trait. And in fact, *human* memory and the capacity to forget are supported by rich mechanisms and are anything *but* simple and well-behaved enough to support mathematical stability. Rather, forgetting is an example of a fundamentally *general* computational notion. The centrality of forgetting to human cognition may instead be a *result* of its centrality to cognition and computation at large. Similar reasoning may apply to other such cognitive groundings. For example:

> [I]ntercultural universality is the result of … rooting in fundamental human cognitive processes (our relationship to measurement and to the space of the senses, basic counting and ordering,…) [which] make them [the invariants / concepts] accessible to other cultures. (p. 43)

To the extent that mathematics is *inherently* grounded in our biology, in a way that could have made it different for a different biology, then classical mathematical realism cannot stand. I insist, however, that some fundamental cognitive processes are actually completely general rather than human. Counting, for example, needn't be "human"; it may instead be a fundamentally computational cognitive activity. Parts of mathematics may be accessible not just to other cultures but to other species – or any cognizing beings; shared *universally*. Having some relevant capacities, cognitive or even physically-grounded, may be a *necessary* condition, for some mathematical topic to apply to the creature (e.g., the *ability* to count, or three-dimensional perception), or for it to be able to *grasp* that topic mathematically. But this does not refute an independent, pre-existing mathematical space that manifested capacities come to *track*, through evolution. To repeat his quotes, I take the following to actually support my own view here:

> [T]he conceptual stability of mathematics, its relative simplicity (it can be profound all the while basing itself upon stable and elementary, sometimes quite simple, principles) can provide the connection which we are looking for with the elementary cognitive processes

Going against his own rich view, fundamental to his approach throughout, by which:

> There is no metaphysics of the ineffable, but rather a concrete, material and symbolic reality of phylogenetic, ontogenetic, and cultural complexity of the human being and his language… This takes us far away from "thought as a formal computation."

3.4.6 Against Timeless Realism

Classical non-biological, non-contingent, timeless realism, Longo rejects, then. For example:

> Can we say that that configuration (a property of the [made-up] game) was already there, a billion years ago? …
> [Although numbers transcend Longo's individual existence,] they do not transcend our human, actually animal existence (as counting is a pre-human activity) …
> The real numbers do not exist, in any sense of a plausible ontology, but their constitution is as objective… (p. 30-31)

Overall:

> Set theory has accustomed people to an absolute Newtonian universe where everything is already said, one just has to make it come out with the help of axioms. On the contrary, mathematics is an expanding universe, with no "pre-existing space," to which new categories of objects and transformations are always being added. (p. 81)

Thus, mathematics can have no *mathematical foundation*, *meta-mathematical corral*, nor *generous arena* (p. 10) (accept perhaps evolving ones), even if these were not merely formal but grounded in the interaction with reality.

His project is to understand mathematics in conventional scientific terms, through scientific methodology. This means going far beyond introspection, and deducing conclusions from findings in:

- Mathematical cognition: Animal abilities, the number-sense, "the spatial component of our intuitions of numbers", relations of "intra-cortical synaptic linkages" to our geometric perception, etc.
- General cognitive knowledge, for example regarding gestalt, or our linguistic capacity and its interactions with other faculties.

Abstract mathematical concepts are manifested in various ways, in cognition too. I will now end this review of Longo with two examples for how

> the signification of concepts and proofs relies also, in contrast with a Platonist and a formalistic view, on some particular features of human cognition. These features precede language, resort from action and ground our meaningful gestures. (p. 84)

This is highly reminiscent of Lakoff & Núñez's "embodied mind" approach (section 3.3.2.3). Accordingly, one may wonder here too about the privileging of particular features, both methodologically (how they are chosen) and conceptually (why privilege anything in particular at all, metaphysically). As noted, this is a challenge for cognitive approaches in general.

3.4.7 The Continuous Line without Thickness

The line without thickness is "a constitutive gestalt of mathematics":

> From the cognitive point of view, one can refer, first of all and simultaneously, to the role of:
>
> - the jerk (saccade) which precedes the prey,
> - the vestibular line (the one that helps to memorize and to continue the inertial movement),
> - the visual line (which includes the direction detected and anticipated by the primary cortex)
>
> …
>
> [T]his line without thickness is the pre-conceptual invariant of the mathematical concept, invariant in relation to several active experiences; it is irreducible to only one of them. This invariant is not the concept itself, but it is foundational and is the locus of meaning: we do not understand what is a line, do not manage to conceive of it, to propose it, even in its formal explication, without the perceived gesture, without it being felt, appreciated by the body, through the gesture which was evoked by the first teacher, through his/her drawing on a blackboard… (p. 67)

That we may not be able to *understand* or conceive of something through its formal explication alone, isn't in itself particularly controversial. But Longo thus takes such precedence to imply a *foundational* precedence (which is most natural for a cognitive approach to foundations).

3.4.8 The Mental Number Line

The natural numbers, rich in structure (as that mental, *geometric* number *line*, "a constructed mental image" through which we seem to perceive them), are grounded in actions, in space and in time:

> the concept of integer, reached through language, goes back to space again, since number is an "instruction for action," counting and layout in space: a gesture that organizes mental space, the one of "numeric line" what we all share (Dehaene, 1997)… the concept of number as a generalized iteration or invariant constituted by an iterated gesture in the space of action. On the other hand, the origin of this concept lies also in a temporal iteration (p. 72)… the (spatial) reconstruction of phenomenal time. (p. 75)

This grounding is to circumvent the limitations of the formal dimension:

> [Proof-undecidability] is a problem for the machines but not for humans: in fact, the geometric judgment of well-order has a "finite" nature and is quite effective from the point of view of a numeric line (only a finite initial segment is to be considered, though not necessarily computable: this is the segment which precedes an element of the non-empty subset considered). This line belongs to the human mental spaces of conceptual constructions… (p. 75)

Thus, in general:

> What is lacking in formal mechanisms… is a consequence of this hand gesture which structures space and measures time by using well order. This gesture originates and fixes in action the linguistic construction of mathematics, indeed deduction, and completes its signification. (p. 83)

Part II:

A Representation-based Exploration

4 My Representational Approach

As discussed in Part I:

Platonism accepts the objective (and eternal, etc.) existence of abstract mathematical objects, independent of the mind in every way. *Cognitively naïve* platonism *implicitly* accepts the mathematician's unaccounted-for but presupposed ability to simply grasp that abstract reality as it truly is. This, a cognitively informed approach could never straightforwardly accept, even if keeping to platonism. Whether mathematical reality is independent of the mind or not, the mathematician's mind must somehow refer to and relate to that reality. How mathematicians come to 'acquire' mathematical reality, become acquainted with it, indeed *perceive* it – if there is any sense to using that word here[35] – is undeniably a cognitive matter. And we *cannot presume to be so fortunate* as for cognition's vast hidden layers to simply be irrelevant and its many deceptions mute in this context. Positioning the "perception" of independent mathematical objects on par with the perception of physical ones actually *undermines naïve* platonism, given the vast non-conscious, misleading complexities of perception; its pragmatic, heuristic nature does not suffice if we are to keep to mathematics' absolute, precise nature. Any fantasy (which may still have been reasonable for Gödel at his time) of unquestionably reliable, immediate, direct perception of anything about mathematics – ought to be forsaken.

In this dissertation, I take mathematicians' tendency towards mathematical realism seriously (as opposed to neo-psychologism, in continuation of section 2.6). Doing so may still stop short of platonism, which, taking mathematical reality to be independent of the mind in *every* way, is thus outside the boundaries of a *cognitive* foundational approach. However, this deep metaphysical difference is actually not a main concern in this part, nor is the question of non-psychologistic realism. The basic exploratory starting point could be shared with a cognitively-informed (as opposed to naïve)

[35] The main theoretical alternative would be to keep "perception" to the senses only, and posit a metaphysically different place for objects as within pure reason. This needn't be just a terminological matter. Once we account for objects as within by the cognitive system or represented by it, the level of similarity between sensory perception and the processes involving mathematical objects becomes a scientific one.

platonism (if there could be one) and with neo-psychologism: The mathematical objects and structures are there *for the mathematician*, whatever they are, and whether they exist (independently of the existence of some mind) or not. The mathematician *seems to* perceive and come to form a representation represent of them. This calls for a *representational* foundation (Section 1.5). The mathematician's mental representation is the main gateway to mathematical reality. In the cognitive details that underlie it, there lies the possibility to take the exploration of mathematics beyond introspection. This is the relatively non-controversial starting point, as metaphysically neutral as possible. Hopefully it may, later, lead to the right, full-scale metaphysics, or at least rule out some non-viable options. For example, through the details of the form of representation, it may become clear whether these represented objects are also *created* by the mind and in what sense. But for the rest of this dissertation, I will set aside the grand metaphysics (as much as possible) and attempt to demonstrate how even a first, very partial cognitive-scientific consideration of the mental representation of mathematical ontology could provide insight into it. My approach is methodological, not metaphysical. Other approaches may complement it rather than contradict it.

4.1 Preliminaries to Representation

Before addressing concrete cognitive issues, a simpler preliminary position is to accept that the mental has a computational basis. This means that in particular, all mechanisms that participate in the activities of reference to, and representation of, mathematical ontology, are computationally constrained. A bit more controversially, I will take computation as Turing bound. This is still a rather standard assumption in general cognitive science, with theoretical considerations to support it (Chalmers, 1994). (It is also likely a suitable way to start, before running into challenges and issues with the background computational assumption, which may then require a reevaluation).

One aspect of the assumption that is problematic is with regard to asynchronous computation, which may be fundamental to the brain's operation, and particularly how it reacts to the environment. But this is not relevant to the cognitive underpinning of mathematics (which I take to be unembodied and not inherently involving response to the external world), at least by default.

Another more relevant assumption here is the finiteness of the computation (allowing at most for the idealization of an infinite discrete tape). A stronger infinitary computational basis could provide *drastically* different support of infinitary mathematics in particular, including even for ℕ, bypassing the challenges of logicism (as in Section 1.1). I adhere to the fairly standard finitistic assumption for computational cognition without further justification. This assumption, in fact, *makes infinity (in its varied forms) part of what a cognitive foundation ultimately needs to explain and account for*.[36] I will term the implication for what representations could be *essential finitism*. It does not entail a directly mathematical, strict finitism. For the minimal example, ℕ may well be a valid cognitively-represented structure or even object (e.g., as a set). Recall (Section 1.6) that the cognitive foundation need not simply be equated with the mathematical objects themselves (i.e. one mental object per mathematical one); it merely needs to bring the objects about, somehow. However, essential finitism does mean that ℕ's conception cannot be implemented "literally" as a collection of all of the natural numbers; its representational underpinnings must be different. Essential finitism is thus posed as a critical challenge to a cognitive foundation of mathematics (to which I relate along the way, but do not purport to offer a solution, unlike, for example, Lakoff and Núñez, Section 3.3.1).

4.2 Cognition in the Abstract

The philosophical-scientific project is holistic. Ultimately, we would have a full cognitive account of the workings of the mathematician's mind with respect to the mathematics to which she attends, and the philosophical picture would have to reconcile with that. Cognitive science in its current state is not at that point. The question, methodologically, then, is how to proceed using what the science *does* already have to offer (which may simply be insufficient to account for fundamental aspects of mathematics).

The natural, "privileged" path, as discussed in Chapter 3, is to turn cognitive science's standard disciplinary tools and research methodologies towards mathematical topics.

[36] Benacerraf and Putnam question that it should "be deemed desirable to avoid reference to the infinite" in foundations. They say that "in point of fact, it is very hard to find reasoned and even moderately detailed argument on this point" (Benacerraf, et al., 1984 p. 6). Although I, too, do not provide a moderately detailed argument here, I suggest that the standard, finitistic view of modern cognitive science (and the unfolding of what finite computation can achieve) further supports this common foundational sentiment which they oppose.

This has brought about mathematical cognition as a field, producing various cognitively-interesting results (with tricky philosophical implications). This approach, however, is effective mostly for the directly applicable, "real-life" mathematical topics (such as numbers or geometry). The entire edifice of modern mathematics is deeply layered in its abstraction, and thus current cognitive science can say very little of it in a *direct* fashion, that accords with the science's own standards of research. How could cognitive science embed within its general framework the psychology behind, for example, the student's grasp of naïve set theory? (The topic of Chapter 8). It could approach it directly, empirically – through behavior of students (with limited room for their subjective reports), in combination with brain imaging. But their objective behavior is not informative enough for our purposes (especially in its narrow sense exhausted by the familiar formalized logical language). Error patterns, learning curves, choices between multiple valid mathematical paths, and the like could help differentiate among suggested models, but they cannot come close enough to telling the complete story. As for imaging, and neuroscience at large, at its current state, it cannot produce a detailed-enough account of such high-level phenomena.

One choice would be to simply wait until the science gets there. But modern cognitive science already offers more than that. Along with the particular findings that are accumulating, it also produces generalized conceptualizations that are seen as central to the workings of the mind at large. Confronted with a particular cognitive question of which the cognitive scientist, and perhaps cognitive science, know nothing, she may nonetheless have an educated guess of what will happen (e.g. behaviorally) and perhaps what is happening underneath (e.g., what the fMRI will show). In this, there is much *implicit* substance, along with a *theory* of cognitive science (that makes some of it explicit). To bring in but one prominent philosopher of cognitive science, Andy Clark:

> In a 2012 paper the AI pioneer Patrick Winston wrote about the puzzling architecture of the brain – an architecture in which "Everything is all mixed up, with information flowing bottom to top and top to bottom and sideways too." Adding that "It is a strange architecture about which we are nearly clueless".

It is a strange architecture indeed. But that state of clueless-ness is mostly past.[37]

The alternative (to fine-tuned mathematical cognition research), then, would be to try to *harness* this general understanding and its far wider explanatory reach; to consider mathematical thinking (with regard also to the potential philosophical implications) through *general cognitive principles*. The central working-hypothesis would thus be as follows:

> *There is no a priori reason to assume that when it comes to mathematics, the mind's workings vary substantially from its general ways.*[38]

This is not to say, of course, that mathematics does not have a special place in the mental realm (and possibly even in the brain, e.g. (Amalric, et al., 2018)). But the finer details of how the general principles are to be applied to mathematics, and what it is about mathematics that justifies which deviations from this general rule, should be handled carefully.[39] In these details, for example, lies a central difference between my approach and that of Lakoff and Núñez (2000), which also pursues a similar *general* strategy (Section 3.3). I reject two of the pillars upon which they base mathematics:

a) Metaphors: I opt instead for a more standard abstract-centered approach.
b) Embodiment: I instead keep mathematics within the confines of information-theoretic cognition; "embrained" (in a sufficiently abstract sense).

Instead, I will:

a) Adopt other common themes. (Chapter 5)
b) Provide an abstract framework that puts our interaction with mathematical objects on par with physical ones. (Chapter 6)

---

[37] From a blog post on his book, *Surfing Uncertainty: Prediction, Action, and the Embodied Mind (Oxford, 2015)*

[38] Its general ways may actually include much mathematical content, even in apparently non-mathematical tasks, as Sloman has highlighted (private communication). Putting it the other way around (as he does), for its general ways, involving "apparently non-mathematical tasks, there is no reason to assume that the tasks really lack mathematical content". (See also (Sloman, 2008; Sloman, 2008)).

[39] The intricacies of general cognition within mathematical cognition are evident in contemporary research in the field too; see e.g. (Hohol, et al., 2017).

c) Construct an analogy between a phenomenon in visual perception and a mathematical one (thus admitting a rather literal reading of mathematical "perception"). (Section 8.3)

It would have been ideal to build solely on existing, well-established theories. But what is known has not yet been put into abstract-enough terms to be applied, ready-made, to mathematics. We thus have to generalize it.

The General Cognitive Principles (GCP) strategy, then, is as follows:

i. Adopt familiar principles, conceptualizations, and theories that are already in existence from cognitive science (particular beliefs or understandings of how the mind and brain function in some respect).
ii. Reframe these notions in abstract-enough terms, so as to generalize over both A) the standard realms already attended to by the science and B) the mathematical topic of interest (particular or foundational).
iii. Adapt and apply these notions to the mathematical topic in order to draw *suggestive* conclusions.

The starting point of GCP, point i, is crucial. Building on a misled conceptualization or theory (of interpreted empirical data) would undermine the support for the conclusion. My own project here is to apply cognitive science *as it currently is* to mathematics and its philosophy, and I attempt to innovate only in the latter. Thus, as a general principle, I strive to build on cognitive principles that are as widely accepted[40] as possible (which is a far-reaching ideal, in today's lively arena of theoretical debate and shifting conceptualizations). This choice is a pragmatic one, which is quite risky. It may certainly happen that some widely accepted conceptualization is nonetheless misled. It may instead be that the mainstream conceptualization should have, but cannot, account for some mathematical issues. In this case, the attempt to base mathematics on actual cognition would be better *reversed* into a demonstration of the insufficiency of the current scientifically accepted beliefs that have been selected for a starting point.[41]

---

[40] This, for (Lakoff, et al., 2000), may be arguable.
[41] An approach suggested and favored by Sloman.

Another risk in GCP-i is in attributing the general empirical findings of cognitive science, which have been gathered with regard to the general population (or at least among the WEIRD[42] societies), to the extremely particular sub-population that is mathematicians. This sub-population may surely have some distinctive cognitive features or tendencies – precisely with respect to the issues that matter for mathematics. The cognitive material on which I will be building is quite general and fundamental, and it seemingly pertains to mathematicians just as much. This would require further scrutiny, in terms of *targeted* empirical research (on these general, non-mathematical issues).

The second of the stages of GCP is the most substantial and sensitive. GCP-ii is an 'extrapolation' of the theorizing already conducted in cognitive science (the topics in higher mathematics aside), but without the brigades of scientists and philosophers of cognitive science attending to it. This stresses how difficult and dangerously ambitious such an aim is: even *within* standard cognitive science, implicit general understanding is usually not made into an actual theoretical claim; it mainly serves as inspiration for raising and testing new hypotheses. The science (not philosophy of mind) has the 'luxury' or standard of empirical validation, putting aside (for the time being) what cannot be methodologically explored in that way. As cognitive science's ability to approach mathematics is not at that capability yet, we must estimate and take a chance. The standard must be weakened: What I present (as do e.g. Lakoff & Núñez) is a suggestive story of the familiar generalized principles and what, cognitively, thus supposedly occurs behind the mathematics. This story, to theoretically inclined cognitive scientists, should at least seem *plausible* (i.e., quite possible and hopefully even likely); a best guess. This plausibility would suggest that it has passed through reasonable generalization, and squares with the scientists' cognitive informedness at large.

Lakoff and Núñez, recall, downplay the uniqueness of mathematics, taking it (neo-psychologistically) as 'just a part of the mental realm'. As part of this, they trivialize the transition through CGP-ii and skipped directly from CGP-i to CGP-iii. For us here,

---

[42] Western, Educated, Industrialized, Rich, and Democratic. (Henrich, et al., 2010)

however, much of the substance is actually in the transition through the *abstraction* of the conceptualization. This is what determines if a such specific approach is valid and could, in fact, provide true insight for mathematics (prior to the empirical validation). Evidence to support the transition to the abstracted cognitive conceptualization could actually be found in cognition itself: It may be that the principles and conceptualizations on which cognitive science converges actually *are* of a more abstract (or 'principled') nature than the science can currently explore; they may be grounded in general *architectural* considerations (that science has yet to reveal as such, perhaps). CGP-ii is a proper extension of theoretical cognitive science.

Last, there is GCP-iii. The *application* of the abstracted cognitive principles requires a seamless stitching to the mathematical substance, valid from the mathematical point of view, as part of the unified cognitive-mathematical one. (This requirement is not always the case for Lakoff and Núñez (2000), the only reviewed representatives who try to innovate mathematically upon a cognitive consideration, enough to risk it; see Madden (2001); Goldin (2001); Voorhees (2004)).

4.3 Dependence on the Particulars of Cognition

Harnessing whatever cognitive science at large can provide could allow us to make some limited but tangible progress in the cognitive understanding of the subject matter. However: The dependence on the *progressive* nature of cognitive science and its theory means that the most we could hope for from such an account of what is actually going on, is to be a best guess. This dependence entails a provisional, scientific progress-dependent nature for the more concrete philosophical conclusion we may try to draw. Rather than producing philosophical arguments and approaches that may potentially stand on their own, this kind of work thus becomes integrated into the scientific process, with alterations expected to arise from scientific progress. The more permanent value is in *establishing* the integration of cognitive science with the philosophy of mathematics – through its beginning. Thus, philosophically, the general form of the argument is that "under this scientific conceptualization, if it is indeed correct, these conclusions follow…".

Integrating in an essential way particular, concrete topics from empirical cognitive science, may seem to render such an approach neo-psychologistic, accounting for *human* mathematics. However, it is the *methodology* that is human centered, not necessarily the insight it can provide. The methodology reveals the dependence of how the mathematics *appears to the mathematicians* on the particulars of the human cognitive system, and through this, helps uncover mathematical reality as it really is (i.e., less contingent and possibly non-contingent on their makeup). This process can be seen as the further '*dehumanization*' of mathematics, to the point where a mathematical realist (and possibly even a hardcore but cognitively informed platonist, as long as the conclusions are ontological but not metaphysical) could be satisfied. That is the point of *ontological rigor* (Section 2.3), of a completely detailed cognitive account[43].

With regard to the foundational goal of *elucidation* (p. 10), a cognitive account complements the progress *within* the mathematical practice (as served in particular by set theory, but also by the general refinements of concepts). It certainly may produce a call for some revisions in mathematics itself (its ontology and principles, not just its philosophy). However, the cognitive approach at its current infancy may do better to resist going against mathematical practice. I instead *rely* on the historical progress within mathematics. In particular, for the specific points I pursue, I rely on Cantor's development of numbers into the infinite (Chapter 7); and on naïve set theory being formalized through the principle of the limitation of size (Chapter 8). The benefit of hindsight is necessary to start the project and assess its viability. With sufficient demonstration of its explanatory force, a cognitive approach could then begin to "advise" current mathematics too.

The focus is on the informal, partly implicit, evolving way in which the mathematician, as a child and then later on as a student, becomes acquainted with the subject matter's ontology. I demonstrate how she is naturally led to not discern inherent distinctions, and even err about the ontology fundamentally, in particular, systematic, familiar ways. This focus thus sets us apart from the main bulk of works on the psychology of mathematics that exist – in mathematical education. Mathematical education is concerned with *challenges* to competence; the distance from mathematics as we take it to be. The empirical data that is relevant there mainly concerns non-universal

---

[43] Contrast with Lakoff & Núñez, section 3.3.2.1.

differences among humans; with distortions that obstruct the *possibility* of connecting to mathematics, of creating valid reference to its ontology. It is the more *universal* principles and mechanisms, those which underlie the officially *correct* grasp of mathematics (according to the practice), that have a better chance to bear on the philosophy of (universal) mathematics itself. Our focus thus blurs the traditional line between mathematics and mathematical thinking, even if we strive – as I do – to refrain from metaphysics.

4.4 Set Theory

Mathematics, even if embedded within the realm of the mental in some sense, seems to be importantly unique. This suggests some *essence* for that which is mathematical, which a cognitive foundational approach can hopefully capture. This may be hard to achieve. For example, what is it that brings together such distinct mental realms as arithmetic and (pre-discretized, e.g. Euclidean) geometry? They both allow for proofs, seemingly absolutely or maximally stable (à la Longo). But is this stability perceived after the fact (i.e., mathematics is what happens to be stable enough; Section 3.4.4), or is it of essence? Why these topics? Could we say in advance, for a different mental realm (e.g. emotions?) whether some distinct, novel part of mathematics could originate from there? There may be some deeper cognitive underpinning that determines the *potential* for stability and proofs. Through it, it would be possible to demarcate the scope of mathematics. However, even so, a cognitive approach to mathematics need not amount to a *uniform* (rather than just unified) account of mathematics at large.

Mathematical practice comes to carve mathematical reality into fields, and the manner of this carving tracks and reflects important differences in their base ontology, the interaction of concepts within and between the subjects, the domains of applicability and their level of abstractness, and the aesthetic intuitions and required mental capabilities that influence a mathematician to prefer or excel at some fields rather than others. Each and every mathematical field could benefit in some way from its own cognitive analysis. Furthermore, appropriate analyses may rely on different tools and concepts to account for what is going on, for instance, relating geometry, separately from arithmetic, to the brain's vast support handling the spatial (Dehaene, et al., 2011) or relating topology to some other cognitive modules beyond those required for rigid geometry.

Ultimately, the various cognitive expeditions into the various mathematical topics would be integrated into the overarching account that, in reality, is unified (though not necessarily uniform). This is the true *generous arena* (p. 10). Achieving this integration may be a *necessary* step on the way to a metaphysical account of mathematics at large.

This hypothetical discussion aside, the point I would like to make is simpler: Bringing together cognitive science and the philosophy of mathematics is an enormous task. And revolutionizing the foundations of mathematics, as arguably called for by modern cognitive science, and as the approaches covered in Chapter 3 try to do (Brouwer and Lakoff & Núñez) or at least suggest (Dehaene and Longo), is certainly a monumental undertaking. Philosophically, these tasks are rich in detail and extremely delicate. Merely relying on *any* part of cognitive science to provide *any* insight on some *particular*, restricted familiar topic in the philosophy of mathematics – if done properly – would be novel enough. This course is a legitimate, if perhaps not the only, way to start.

The rest of the dissertation thus revolves around two very particular points that I aim to make. The central tool that I develop along the way is much more general, and the methodology to be commenced (as described in this chapter) is certainly so. However, I do not pursue the large-scale foundational project directly, in the general terms that have dominated these introductory chapters. The particular points are both within set theory. Any mathematical field could likely serve for beginning the cognitive project through demonstration, but set theory is privileged in two senses of interest:

<u>Infinity</u>

Potential infinity already poses challenges for any cognitive approach (as for any other foundation to which we relate). Actual infinity may take these challenges further, but the non-denumerable, higher infinities, some of them so integrated into (and seemingly crucial to) modern mathematics, take the difficulties to a whole new level. A cognitive basis can no longer be object-complete. The mathematician presumably grasps and reasons about a reality of objects most of which transcend her grasp. Accordingly, non-denumerability has been a center of controversy between the classical "cognitive" approaches (Section 3.1) and "classical" mathematics. Although a cognitive approach may not have to abide by those same phenomenological principles, it still must rise to the challenge, somehow. In particular, object-incompleteness suggests a more intricate account than that of mathematical objects as mental objects. Thus, although I do not offer a solution to these difficult matters, set theory has the value of bringing in this central tension.

Foundations

For philosophy, set theory is not just another mathematical field (Chapter 1):

> It's more or less standard orthodoxy these days that set theory - ZFC, extended by large cardinals -- provides a foundation for classical mathematics. (Maddy, 2017)

According to this orthodoxy, providing the foundation is done by grounding mathematical ontology and epistemology (though not metaphysics) in sets and their theory:

> Set theory is the foundation of mathematics. All mathematical concepts are defined in terms of the primitive notions of set and membership. In axiomatic set theory we formulate … axioms about these primitive notions … From such axioms, all known mathematics may be derived. (Kunen (1980), in (Maddy, 2017))

Beyond being sufficiently rich for grounding ontology and epistemology, a foundation should also be theoretically simple (following p. 14). (The promise of) being able to manage with a single theory of a single ontological type (sets) understandably seemed (a century ago) to fare quite well with the vague aspiration of simplicity. Moreover, any full foundation other than set theory would, in particular, have to ground set theory itself, and the only apparent option, via logic à la Frege, turned out to be unattainable. However, from the modern cognitive point of view, we know the following:

i. This simplicity may be misleading and may conceal a complexity of which we are unaware.
ii. A computational basis that grounds set theory or at least what it means to us (i.e., its representation), *should* in principle be possible, somehow.

Through Part II, and with (ii) as a working assumption, I will demonstrate (i); namely, that set theory itself – as viewed from its cognitive representation or grounding – is not so 'simple'. It is prone to foundationally-illegitimate cognitive influence; further ontological rigor is required of a foundational basis. Hence, my focus on set theory is not for its foundational value, but quite the opposite, *in the face of* its traditional centrality, which is thus challenged.[44]

[44] This cognitive critique is different from that of the main modern foundational alternative of category theory and homotopy type theory. However, the two share some themes and may well be compatible (Section 1.4).

5 The Architecture of Cognition

For the sake of the "non-cognitive" philosopher of mathematics, I now introduce in (extreme) brief a series of familiar topics from cognitive science to which I make recurring reference throughout the dissertation. These topics are widely general in the workings of cognition. Thus, assuming a priori, for the sake of the GCP strategy (p. 119), that they pertain to the mathematics-related workings too, should be a relatively non-controversial undertaking. The conceptualizations, however, are still under debate and changing all the time. I will select particular introductory sources and leave these complexities for future developments.

5.1 Working Memory

Working memory (WM) is a central theoretical concept of cognitive science that "attempts to address the oversimplification of short-term memory by introducing the role of information manipulation" (Chai, et al., 2018). It is where the 'live' conscious manipulation of information takes place.

Within it (according to the multicomponent model, the most prominent one), "the central executive functions as the 'control center' that oversees manipulation, recall, and processing of information" (ibid.). "The central executive (which resembles an attentional system) is the most important and versatile component of the working memory system. Every time we engage in any complex cognitive activity (e.g., reading a text; solving a problem; carrying out two tasks at the same time), we make considerable use of the central executive" (Eysenck, et al., 2010 p. 217).

The system is central and general, applying to all (or most) issues of (high-level) cognition, particularly to symbolic processing and the empirically explored topics in mathematical cognition too. Given this, it is a relatively conservative generalization to assume (a priori) that the system governs the rest of the cognition of mathematics as well. This generalization's reality would mean that the *limitations* of working memory apply to that cognition just as well.

A central limitation is that the central executive, at the top level of consciousness, is usually[45] considered to hold only three to four items at once. That is, only three to four objects can each be attended to in an individuated fashion, represented as such. The rest of the richness, then, is in what the objects are taken to be (e.g., "chunking" strings of numbers according to some attributed meaning). With respect to mathematics, too, Feigenson provides the following: "I suggest that the hallmark capacity limitation of WM is accompanied by impressive representational flexibility. Only three or four items can be maintained in WM at once; but critically, different types of entities can function as items. These include individual objects, sets of objects, and ensembles" (2011).

5.2 Automaticity

A central distinction that courses throughout many areas in the cognitive sciences is *automatic* versus *controlled* processes (Moors, et al., 2006; Schneider, et al., 2003). One familiar example is the *Stroop effect*, where identifying a word's color – a controlled process – is hindered by an automatic recognition of the word and the color it names. There is no single agreed-upon definition for automaticity (as contrasted with being controlled), and many of its characteristics do not always co-occur, but the traditional view suggests "that automaticity should be diagnosed by looking at the presence of features such as unintentional, uncontrolled/uncontrollable, goal independent, autonomous, purely stimulus driven, unconscious, efficient, and fast" (Moors, et al.). Another important feature is that automatic processes do not reduce the capacity for performing other tasks.

A related also-central distinction that should be mentioned is between *declarative* and *procedural* knowledge (e.g., knowing people's names vs. knowing how to ride a bike): "Declarative knowledge is stored in chunks or small packets of knowledge, and is consciously available. It can be used across a wide range of situations.. In contrast, it is often not possible to gain conscious access to procedural knowledge, which is used automatically whenever a production rule matches the current contents of working memory" (Eysenck, et al., 2005 p. 456). In a nutshell, whereas logic-based approaches keep – by fiat – to declarative knowledge, a cognitive approach to mathematics must

[45] There are, however, other theories for how to interpret its discovered limitations. See e.g. (van den Berg, et al., 2014)

seriously address the possibility that there is some procedural knowledge in mathematics that is not readily available consciously.

Rather than a clear, fixed dichotomy, automated versus controlled processes are ideal points on a dynamic spectrum. In particular, automation through practice (e.g., driving a car, which becomes much less attention consuming with experience) is a central component in the development of *expertise*. "Skill acquisition typically involves.. a progressive shift from the use of declarative knowledge to that of procedural knowledge, and an increase in automaticity (Eysenck, et al., 2005 p. 456). And again, there is no a priori reason to assume that mathematics as expertise would be different.

Overall, automatic processes are extremely important for our cognitive functioning, as they take the load off of our extremely limited conscious resources, particularly our working memory. However, it is precisely this power of theirs that obscures conceptual foundations, as it means that common, central complexities of our mental inner-workings are beyond awareness, hidden from our conscious sight.

5.3 Perceptual Hierarchy

Objects as represented by the perceptual system are highly structured. This structure is largely hierarchical. Many levels of computation underlie the high-level objects as perceived. Here is a standard presentation of the topic:

> Each transduction element [the outer starting point of sensory cognition] transfers information about a small, local aspect of the total energy flow… Yet, our brain's interpretations of the external world, expressed in our perceptual experiences, are immediately holistic and ecologically meaningful. We see and hear objects rather than their local constituents. Thus, we immediately perceive a complex sound as that of a breaking glass, or of a spoken word, and automatically identify the source that emitted it. On the other hand, we have no immediate percept of the frequencies that compose these complex stimuli. Similarly, we perceive faces and houses, but do not immediately perceive the retinal position of their parts. As summarized by the Gestalt psychologists, looking outside the window, we see a forest rather than the trees composing it.
>
> It is now largely agreed that the gap between local sensation at the peripheral sense organs and global perception is mediated by local-to-global processing hierarchies. Although the pattern of connectivity is not strictly hierarchical…, its general nature,

expressed at the anatomical, physiological and recently also in imaging data, reflects gradually more global representations in areas farther from the periphery, even at the single cell level. Thus, in the visual modality, which has been most intensively studied, lower level representations (beginning at the retina) are believed to extract basic, general purpose local primitives, such as oriented bars or basic colours. At the other extreme, high-level representations are closely related to our global percepts… The transition from low levels to high levels, i.e. from local to global, is [generally] gradual…

[P]hysically different stimuli, whose low-level representations are very different, may belong to either the same external object or similar ones, which belong to the same high-level category. This convergence creates an essential generalization (e.g. a single object is perceived as same whether near or far), but at the cost of physical resolution. (Ahissar, et al., 2008)

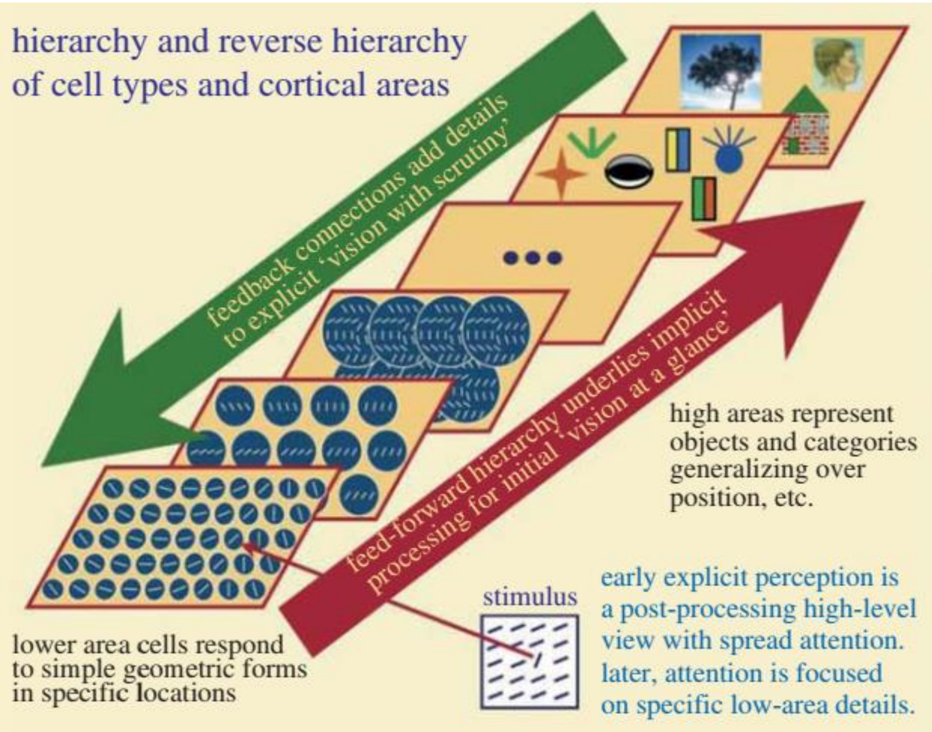

Fig. 3 Schematic of the local-to-global processing hierarchy. Source: (Ahissar, et al., 2008)

Higher levels (the "cortical areas" above) bring together many lower-level computations, constructing deeper, more complex representations. Growing from local to global responses, it must be noted, is just one expression of this phenomenon (if one that is easy to explore).

While the computational structure underlying the perceived objects is relatively readily available for neurological exploration (as has been going on for decades), the phenomenology of it all is less agreed upon. But to bring in one such relevant theory:

> The reverse hierarchy theory is a concept that aims to link between the hierarchies of processing and the dynamics of perception… It proposes that, by default, rapid perception is based on high-level representations alone…
> [I]t implies that our typical perceptual experiences (i.e. our conscious perception) reflect *only* the information stored at higher levels. Thus, if high levels are global, abstract and represent the 'gist' of ecologically relevant objects and events, it is only this gist that will be immediately experienced… when a specific [neuron] population denoting an object category (e.g. car) is activated we can immediately tag it as a car even though we have no immediate access to the details of its spatial components. From an ecological perspective, this limitation is a by-product of the need to generalize across different instances of the same object or even of similar objects. (Ahissar, et al., 2008)

Worldly objects are *not* the pre-given basis of the universe as *cognitively represented*, but the result of a processes-based construction. The choice of how to carve perceived reality into objects – including how many, how high the level, what to abstract away from, what to approximate, and what to attend to – is a matter of computational and functional optimization, not a part of physical reality (although certainly highly dependent upon it too). This described hierarchical structuring of cognition, rather than an issue specific to vision, is quite general. In this way, the architecture manages with an extremely limited working memory that supports only a few conscious objects (which means they must be high level, and approximate).

5.4 Language

Within this picture of the hierarchical complexity of objects and their processing, and our phenomenology of it, through its top – where does the linguistic representation of it all fit in? Very generally, language is related to that upper-most level. This is not to belittle the vast cognitive complexities that are to be found in the linguistic processing itself. However, the words in our possession for designating objects (e.g. "house") are for what we perceive and handle consciously, namely the high-level objects (e.g. house).

Another important point to note it that language is constructed for efficiency. Thus, in particular, there is no point in being able to express explicitly what is the same for everyone, whether because it is innate, automated, or just obvious.

6 Cognitive Object-Constitution

The interface between reality and the cognitive system is a low-level one. Micro-elements of reality (e.g. photons) that reflect what is going on in it interact with the computational atoms of cognition (e.g. photoreceptors). Yet, reality seems to have (at least approximately) a higher-level structure of *objects*, from rocks to chairs to fruit to people. Metaphysics aside, such conglomerates, as high-level constructs, seem to play a central role in the workings of our minds; they are the very fabric of the conscious mental realm, which does not care for individual micro-percepts nor atoms. Even if the cognitive system's ultimate function is to perceive, predict, and (through motor control) manipulate streams of bits, high-level objects, it seems, are crucially instrumental.

The ontological structure, the various types of objects, are not simply "given" to the system. Whatever the real-world objects are, the system has no direct access to them *as such*. It can only do its best to reconstruct them by retracing the contours of reality itself, *with respect to the influence the objects have on the cognitive processes*. This includes every aspect of the objects that interacts with the system – i.e., both *perception* and *action*. The system needs to "understand" the ways the objects influence it (an apple on the table could ease hunger), the ways it can influence them (bring apple to mouth), and how they influence other objects which influence the system indirectly (the dog is nearer to the apple, and will probably take it first). The system thus needs to *construct* its object-*re*presentations, what it takes the real-world objects to be; what they mean for it.[46] To borrow the words of Fodor and Pylyshyn: "What you see when you see a thing depends upon what the thing you see is. But what you see the thing as depends upon what you know about what you are seeing". (1981)

[46] An alternative approach is that high-level objects never do play a part as such in the functioning of the cognitive system. There is perception, data processing, prediction, and action, but nothing is ever about objects. This can at the very least find support in modern machine learning, where data prediction and compression is all that there is, adjusting weights in the (idealized, computational) "neural network".
But to the extent that objects do play a conscious part in mental functioning, this should still be accounted for, even if in terms of or deduced somehow from the computation that takes place (e.g. be extracted somehow, even if we do not know how yet, from the highest levels of a deep neural network). The conscious part is part of what needs to be accounted for. And as mathematics is explicitly completely committed to objects, let us for this project take it that there are such things, mathematically and cognitively represented.

Figuring out the real-world structure, learning the types of objects, determining their representations, takes place at two stages. The first is the whole evolutionary path, which may hard-code some fundamental aspects into the circuitry; innate knowledge (e.g., “these are faces, and they are worth looking at”). The rest is what is underdetermined and learned by the system itself[47] (through learning processes optimized by evolution). The actual division of labor between the two stages is an extremely complicated and very much open issue. An issue on which, in regard to our own specific *mathematical* interests, I take a stance in Section 6.7.

This general process, of converging onto the types of objects (which are not pre-given), of figuring out the ontology of a realm, is what I shall term *objectification*, or more elaborately, *cognitive object-constitution*. It is what should probably be supported innately (though it in itself may to some extent be altered or optimized according to outside data). A "Kantian preliminary" if you will, part of the system itself.[48] By “the objectification system” I refer to the part of the cognitive system that handles this directly (making use of various other systems, and providing input for others). It need not be implemented as an actual, separated cognitive module (neurologically or computationally), but for our conceptual purposes, it may be thought of as such a sub-system.

Following some background, I shall put forth a fundamental general framework (and abstract conceptualization) for objectification – in whatever realm that may be. It is not restricted in advance to any sub-realm of cognitive phenomena. In whichever spaces “objects,” of whatever kind, seem to us to make sense and be involved, I strive for the framework to reign over them: physical inanimate objects, living ones (to whom we may wish to attribute more than just their body), and more – and mathematical ones as well. The aspects that I stress, however, are limited to what is required for supporting what I later build upon the framework. In particular, for objectification in general, a far-reaching generalization would be required, for handling *probabilities*. This and other

[47] The system's operation may also bring it new knowledge regardless of any outside influences. But as long as there are no alternatives for which the choice between them is influenced by alternative environments, it still falls under the first type; it is still innate.

[48] It may be, however, that even such a fundamental issue is one that the system *learns* instead, converging on an object-oriented representation of reality as the solution to some more fundamental principle yet, e.g. optimal prediction.

central aspects that lack the in-depth exploration they deserve are mentioned as the work progresses for the sake of the larger picture of objectification. What I strive for is only that this account is reasonable enough as a sketch of such a system. A proper account far exceeds the scope and goals of my project.

The framework itself is a most abstract one. Its aim is to suggest and clarify (if not determine completely) the conceptual space in which the task of objectification takes place. It provides the terms in which actual such systems should be understood and explored, be they natural or human-engineered ones. It is the mold that their many implementation-related details should fit into.

6.1 Background Areas

The objectification framework I offer, as an extremely fundamental one, relates to many issues in various areas. To contextualize it, I review in this section some of the most relevant areas. Building the conceptualization of the topic from the ground up necessitates "reinventing the wheel" to a considerable extent. However, there does not seem to be on offer an abstract-enough, unified framework to satisfy our needs, for relating, from both the formal/engineering and the empirical/biological perspectives, the handling of mathematical objects to that of non-mathematical ones.

In continuation of the generalization challenge of GCP-ii (p. 120): These formal/engineering and empirical/biological sides may draw inspiration from one another but keep to their own goals, methodologies, and standards. My philosophically minded account does not need to stand up to any side's own innovative research standards, but it does need to transcend mere inspiration, to actually seem reasonable from both *informed* perspectives. Some examples and concepts from the areas covered in this section are woven in, as the framework unravels, to stress the relations for the sake of a more integrative view. After presenting the framework in full, in Section 6.6, I demonstrate how it can be reassessed with respect to cognition.

6.1.1 Model Theory

Model theory formalizes the basic notion of a *type* (of object) (Hodges, 1997). Examples of types (in arithmetic, say) are as follows:

- Being an even number (– a type of number)

- Being larger than 3 and smaller than 5 and all other facts about 4 (a so-called complete type, having all facts about it determined)
- Being larger than n – for each n ∈ ℕ (a type that no natural number satisfies)

Generally, a type is conceptualized as the set of formulas (with a free variable) that hold or could hold of an element (and more generally between an n-tuple of elements). This is in line with the general spirit of logical foundations, of approaching mathematical objects and structures through the sentences that are true of them (Section 1.1). My framework offers and builds on an analog, in the different, procedure-grounded spirit and focus of this dissertation.

6.1.2 Object-Oriented Programing

Programing has traditionally been focused on the various data-processing functions (calling more functions in turn). However, a central modern programing paradigm rearranges the construction of programs to revolve around *objects*. These are constructs to which both data elements and functions ("methods") are then attributed. (The paradigm is appropriately called object-oriented programing; OOP).

This move shares much with the framework I present. But very much unlike mine, it has been deeply explored, in terms of possible alternatives, semantics, etc., offering a rich formal apparatus. However, OOP aims to serve the designer or programmer, to ease the *transference* of the ontology that is already in her head into working code, or lead her to clarify it sufficiently for that. My framework, instead, is focused how these constructs could *come to be*, in the first place, come to be learned *automatically* (i.e., with no implicit further programmer intelligence assumed). This difference has many design implications. For example, OOP sanctifies a simplifying separateness ("encapsulation"), whereas cognitive objectification must focus on the rich interactions between the constructs and how they come to be acquired. Thus, there are some important differences in how I shall approach the topic.

6.1.3 A.I. & Robotics

Going beyond the explicitly programmed ontological structure into learning from reality about it, autonomously, is part of the purview of the vast field of Machine Learning and Artificial Intelligence. For just one particular example of an exploration that is related to ours:

Some research in machine learning attends to the integration of different modalities (with respect to perception, at least)[49]. It builds on the vast data and modeling experience gained in the different modalities (which makes sense for the state-of-the-art performance). But there has yet (to the best of my knowledge) to be a unified approach that integrates the various inputs and outputs from the ground up, independent of modality. Since for mathematics an abstract framework is required, independent of the particulars of some physical modality (if there could be one), this would not suffice.

Beyond general machine learning, our issue is with objects (i.e. outside of the system) and how the system *interacts* with them (rather than just perceives and, say, classifies them). And embodiment – is the purview of robotics.

To be sure, being grounded in the physical (and in spatial geometry) entails many differences from our mathematics-oriented purposes, and makes the generalization a deep challenge (the challenge of this chapter, and a central one for my approach as discussed in Section 4.2). Much can be achieved through more concrete frameworks, in which much of the setting is "coded in" rather than designed to be discovered. Though, even within the confines of their physical, spatial goals, robotics has just begun to attend to deeper background principles, having the robot learn for itself substantial parts of the structure of the ontology it is designed to handle. This is the subfield of "developmental robotics, which aims to show how a robot can start with the 'blooming, buzzing confusion' of low[-]level sensorimotor interaction, and can learn higher-level symbolic structures" (Modayil, et al., 2008). More philosophically, the authors have noted, "Instead of considering an object to be a physical entity, we consider an object to be an explanation for some subset of an agent's experience. With this approach, the semantics of an object are intrinsically defined from the agent's sensorimotor experience". This move is in the same direction we take, but it is more concrete, which allows the authors to actually offer early working systems (within their limited scope and goals[50]). My more general framework is not developed this far, and it explores only a subset of

[49] For a recent example: (Aytar, et al., 2017), (Kaiser, et al., 2017).
[50] For example, (van Hoof, et al., 2012), who optimize the robot's exploration informational-theoretically to tell apart objects in the visual scene, conclude the paper with their plan "to use our approach to scene segmentation as a stepping stone for learning object properties. Besides intrinsic properties such as shape and mass, it is important for a robot to learn which actions an object affords and how these actions change the state of the objects in the environment".

the issues in accordance with its philosophical purpose. In particular, it leaves unexplained the basic issue of *recognizing* objects (as their type), and related issues like tracking them. However, the framework must in principle be extendable into a working system that includes those aspects. And unlike for robotics (to repeat), it must go beyond relating to biological systems through mere "inspirations", towards capturing general (if idealized) aspects of the workings of such cognitive systems.

To living cognitive systems we turn next. Before that, I briefly demonstrate how robotics, sharing our direction, could thus shed light on my own endeavor (and on the cognitive topic) too: "experiential knowledge has some significant advantages over knowledge directly encoded by programmers as it intrinsically captures the capabilities and limitations of a robot platform. The advantage of experiential knowledge has been demonstrated in the field of robot mapping where robot generated maps are more effective than human generated maps for robot navigation" (ibid.). The reasons underlying this may be relevant to the living just as much. The particular circumstances into which one is born vary, as do the particulars of the body – and of the brain. Developing one's capabilities intrinsically with respect to the world makes the adjustments to these variabilities inherent. In comparison, direct encoding of the concrete – by programmers or by evolution (i.e. innately) – cannot be as flexible. This suggests that living systems too may be better off open ended, set up with only the most essential pieces innately coded in – and with the general rules for acquiring the rest of the structure of reality.

6.1.4 Cognition

The cognitive revolution gave center stage to representations (section 1.3.2). And representations – are first and foremost of objects. Objects, in turn – first and foremost means physical ones. To put it bluntly, then: "The human development is a very motivating example of efficient learning about the environment without explicit supervision. Indeed, object representation is considered as one of the few [pieces of] core knowledge that form the basis of human cognition (Spelke and Kinzler 2007)" (Lyubova, et al., 2016). The referenced authors, in a related paper, elaborate as follows:

> The core system of object representation centers on a set of principles governing object motion: *cohesion* (objects move as connected and bounded wholes), *continuity* (objects move on connected, unobstructed paths), and *contact* (objects influence each others' motion when and only when they touch)… These principles allow infants of

a variety of species, including humans, to perceive the boundaries and shapes of objects that are visible or partly out of view, and to predict when objects will move and where they will come to rest… research with older infants suggests that a single system underlies infants' object representations. For instance, 5-month-old infants do not have more specific cognitive systems for representing and reasoning about subcategories of objects such as foods, animals, or artifacts… or systems for reasoning about inanimate, non-object entities such as sand piles… Finally, infants are able to represent only a small number of objects at a time (about three…). These findings provide evidence that a single system, with signature limits, underlies infants' reasoning about the inanimate world. Investigators of cognitive processes in human adults have discovered evidence that the same system governs adults' processes of object-directed attention, which accord with the cohesion, continuity, and contact principles and encompass up to three or four separately moving objects at any given time. (Kinzler, et al., 2007)

Just about every piece of this could ideally be reflected within my framework, if it is to account for the representational constitution of physical objects rather than just of the *differences* between various types of physical objects. This is important for generalizing beyond physical objects to mathematical ones too (or else the system might depend on the physical particulars of the core system that supports it). Let us begin then by deconstructing the perception of physical objects itself.

Within a specific modality (paradigmatically, vision), the official 'objects' are reconstructed and recognized through simpler perceived 'objects', such as edges (which are in turn reconstructed, ultimately through single photons that represent much less than a full edge). But they are also reconstructed and recognized through a multitude of particular properties. This is explored under headings such as *binding* or *conjunction* – which is also within our scope: "In the visual domain, if color, motion, depth, and form, are processed independently, where does the unified coherent conscious experience of the visual world come in? This is known as the binding problem and is usually studied entirely within visual processes" (Wiki-MSInt, 2018). "However", the quote continues, "it is clear that the binding problem is central to multisensory perception". At a sufficiently general level, combining the information from the two eyes should not be *essentially* different from combining, for instance, vision and hearing (and likewise at many other levels of the representations of objects). Ultimately, physical objects are cognitively represented through a multitude of modalities. How this comes about is explored in cognitive science under the heading of *multimodal* (or *multisensory*, or *cross-modal*) *integration*. It, too, should be within the scope of my framework, as it is to account for *mathematical* objects through somewhat analogous integrations later on. If multimodal integration were completely innate, then it could not be related to my approach to learning of mathematical objects. However:

> "[N]eurons in a newborn's brain are not capable of multisensory integration, and studies in the midbrain have shown that the development of this process is not predetermined. Rather, its emergence and maturation critically depend on cross-modal experiences that alter the underlying neural circuit" (Stein, et al., 2014)

Beyond the core system of physical objects as just that, "adult humans also have developed more specific knowledge of subdomains of objects such as foods and tools" (Kinzler, et al., 2007). This also falls under the scope of my framework, which is to cover the many issues regarding *types of objects* (and through generalization then covers mathematical object representations too). This brings us to the classic, well-explored topic of *categorization* (see e.g. (Cohen, et al., 2017)).

Categorization is, first and foremost, about how humans (in particular) classify, come to recognize the objects in the world as belonging to certain categories at various levels (where the division is usually taken to be in the world itself). Importantly, (supporting Wittgenstein), categorization breaks away from the Aristotelian classic of necessary and sufficient conditions, for example offering graded categories (wherein a robin is more of a bird than a penguin is). The topic offers much regarding how we represent categories instead, in terms of prototypes or exemplars, for example.

On the other hand, with classification as its main task, classical categorization theory puts the spotlight on perception. Later on in life, we come to learn of new actions that we can *do* to an object, or other facts about it. But cognitively, the object is already there, classified; what it affords is not *inherently* part of its conception. Later on, through mastering finer distinctions, we also come to break apart divisions further (e.g., distinguishing beagles from boxers). But what matters in classification is always how the world's objects are divided, not finer distinctions that still can matter cognitively. (The contrast with my framework is clarified shortly).

The different focus of my framework, on object-representations as constituted upon our interaction with the objects, takes us some way towards Gibson and his “ecological” approach:

> It was generally assumed until about 25 years ago that the central function of visual perception is to allow us to identify or recognise objects in the world around us… Gibson argued that this approach is of limited relevance to visual perception in the real world. In our evolutionary history, vision initially developed to allow our ancestors to respond appropriately to the environment (e.g., killing animals for food; avoiding falling over precipices). Even today, perceptual information is used mainly in the organisation of action, and so perception and action are closely intertwined. (Eysenck, et al., 2010 p. 121)

I take the principled integration of perception and action as fundamental, far beyond for just *physical* objects and their *visual* perception. The precedence of perception over action in early infant life is assumed herein as coincidental, given how some abilities simply develop before others. However, the general principles of cognitive object-constitution, which are to apply to new types learned and mastered at later stages in life too, integrate actions into the very essence of what an object is for us. This fundamentally integrative approach becomes quite risky in relation to the brain, which has rather separated systems, which are sometimes conceptualized as managing perception versus action – the so-called “two-streams hypothesis” (see e.g., (Eysenck, et al., 2010 p. 47)). Attending to these complexities is left for Section 6.6.2, after the framework has been presented.

What is certain is that classifying objects successfully is but one task, if preliminary and necessary. This brings us to *affordances*:

> Gibson (1979) claimed that all potential uses of objects (their *affordances*) are directly perceivable. For example, a ladder “affords” ascent or descent, and a chair “affords” sitting. The notion of affordances was even applied… to postboxes (p. 139): “The postbox . . . affords letter-mailing to a letter-writing human in a community with a postal system. This fact is perceived when the postbox is identified as such.” Most objects give rise to more than one affordance, with the particular affordance influencing behavior depending on the perceiver’s current psychological state. Thus,

> an orange can have the affordance of edibility to a hungry person but a projectile to an angry one. (Eysenck, et al., 2010 p. 124)

For us here, affordances are integrated into the object-representation itself, and so they are "directly perceived" of the object (including perception of properties required for action, such as an apple's size for grabbing it). Affordances of the object itself can thus be learned into its (extended) representation. Representing objects and perceiving them as their type, hence, is not trivialized. I accept that "Gibson's argument that we do not need to assume the existence of internal representations… to understand perception [and action] is seriously flawed" (Eysenck, et al., 2010 p. 125). The representational theory I offer is more along the lines of what Nanay has termed 'pragmatic representations,' "that mediate between ~~sensory~~ input and ~~motor~~ output" (my strikeout) (2013 p. 3). (The similarity is also in that they need not "have a syntactically articulated propositional structure", i.e. à la logicism).

Gibson's ecological approach also grew into the important alternative within cognitive science that *Embedded Cognition* (Wilson, et al., 2017) is today. It is important to stress that I do *not* go that way. Actions (and perceptions) are considered as fundamentally *cognitive*, however they may come to express themselves at the periphery of the cognitive system (i.e., the body) and from there on into the world. This is particularly important if we are to generalize beyond physical objects, to mathematical ones too. Whatever they are, our access to them is not through the standard senses, muscles, etc. (Contra (Lakoff, et al., 2000), Section 3.3).

The framework is intended to bring together some underlying principles regarding the very task of objectification. But abstract principles can only go so far in bringing together such a wide variety of phenomena at various levels of complexity. Innate adjustments (which are substantial for core issues), the particular specializations of parts within the brain (e.g., those two streams, ventral vs. dorsal), all put limits on what can be explained this way, without concerning the actual empirical details. But the few concepts and principles that we can explore do bear on these topics, which thus provide both challenges and things to build on and use to enrich the framework. For example:

> Part of an agent's development are early failures in the perceptual system. One commonly discussed phenomenon in child development is the lack of object

> permanence, where the child fails to maintain a representation for an object when it is not perceived in the observation stream. Another type of failure is the failure to correctly track objects…
>
> [And also:] Children develop from using motion as an indicator of object unity to using other cues. (Modayil, et al., 2008).

Such issues need not be directly relevant to a framework that purports to extend to non-physical, mathematical objects too. But we can (and I do) reason and speculate as to some more abstract principle that may underlie these cognitive-physical manifestations. The difficulty in this sort of challenge cannot be overstated. For example, merely for multisensory integration, Ohshiro et al. explain the following:

> Responses of neurons that integrate multiple sensory inputs are traditionally characterized in terms of a set of empirical principles. However, a simple computational framework that accounts for these empirical features of multisensory integration has not been established. (2011)

The very abstract outline I shall present, then, could not translate directly into solutions to such problems (before further details yet to be known could hopefully expand it). But the limited scope of what this framework *would* purport to achieve should later on shed some light on mathematics and its objects.

This concludes the survey of relevant background to the framework I now present. A glaring omission is philosophy itself, in which various topics relate to such a framework. The framework is introduced, however, only as a tool, and as such, its proper grounding in the philosophical literature is not crucial (if not beside the point).

6.2 Procedure Arrays

Putting representation before the object, "An object… is a hypothesized entity that accounts for a spatiotemporally coherent cluster of sensory experience." (Modayil, et al., 2008). My similar, abstract version is as follows.

A (cognitive) object-type[51] (e.g., the conception of an apple) is a category of object-*representations*, through which operations on worldly objects are cognitively mediated. It is (cognitively) constituted upon the following:

Procedure Array: a designated set of procedures such that

i. Each procedure operates on the state of the postulated object-to-be, the state of a pair of those (interactions)[52] – or the state of other procedures from the set (indirect operations).
ii. The set (as a whole) is *object-coherent*.

I now unravel the above, starting with the procedures themselves.

The cognitive procedures are the atoms. They are mere computational operations, a priori devoid of any content or intentionality (aboutness). They are *taken* by the objectification system (and by us the theorizers) to operate on *postulated* objects – to query or manipulate them (computationally speaking); i.e., perceive or act upon them (or both). A particular procedure may be part of many arrays. For instance, a hand muscle can be contracted both as part of hitting the ball during a tennis game and as part of slapping someone. But the procedures of an array, together, bring about the object-type – the whole of the possible interactions between the cognitive system and objects of that type. The *cognitive* object-type thus makes for a finer distinction than what is taken to be *types of objects* (an expression I use rarely, only in the following respect): For someone who learns how to notice something or perform a new action with some type of object (like juggling apples), it is customarily still considered to be the same type of object (i.e., in the world). For the framework here, though, we must take such differences (between the old and the new object-type) into account too. I leave open the further issue of *when and why* we come to identify different (cognitive) *object-types* as of the same type.

[51] In Object Oriented Programming, the analogous notion is that of a class.
In philosophy, a related (though different) notion is that of a kind (used mostly through the metaphysically "privileged" ones that are natural kinds). I do not explore this further, and, as usual here, shall keep away from the many philosophical issues lurking around the corner.

[52] I doubt that numerous sets of objects are directly relevant (in working memory, say). Some generalized form of "chunking" probably operates iteratively to bring about some n-way interactions.

(A full account here would need to connect this issue to the actual atoms of cognition, e.g., neurons. For our current purposes, we can do without it and simply assume that enough of these together accumulate to some sufficiently rich and well-defined computational model for the operations on the postulated object-types).

I opt for the broad view that makes room for meta-procedures too. For example, a ball allows for *bouncing* it (i.e., as one of its procedures) by throwing it and catching it repeatedly. Procedures do not have to operate and be about the objects *directly*. Some procedure *P* that is about the object-type *O* may depend on parameters (in memory, as part of the mental state). Another procedure *P'* may query or manipulate those parameters. With all procedures working in tandem, this means that *P'* does query or manipulate O, at least potentially (there is no harm in allowing for vacuous operations, nor operations that may actually happen at a substantially future time). Meta-operations may play the central role in applications of this framework in which we may be interested (and they are probably needed for reflecting the actual neural or computational architecture). For example, the turning of an object requires coordination between the operation of many different muscles (while further sensing, perceiving the change on-line for control). This may be the way to model such an orchestration, and indeed the way to implement it.

And now we arrive at the central requirement, the coordination among the procedures; for the set as a whole to actually constitute an object-type – coherently. Coordination (as a relation between the procedures, across time) is what enriches and indeed brings about object constancy. Turning a physical, rigid object left and turning it right are coordinated (as operations in physical reality) to cancel out each other (in terms of overall effect). Turning it left is coordinated with then perceiving it as having turned left. Without these regularities, part of its physical constancy would collapse; the object-type would lose some of its properties (now ill-defined). If turning it left and then right (by the same amount) resulted in variable and unpredictable results, the notion of its angle towards us would collapse. (And likewise, if turning it right did not then bring it to be perceived as turned right). This disorientation may conceivably happen with bad brain-wiring. But it may also be due to an irregular environment. Thus, the procedures must be taken *in context*, as embedded in their *standard environment*, not just within their cognitive boundary. Object-coherence is thus a very strong notion of coherence, not just syntactic (concerning just the computations themselves). [53]

I take coordination as an abstract relation here. The implementational manifestation and implications of coordination I leave unspecified. But just to illustrate, consider the coordination between moving an object – and perceiving it as moved. When moving one, perception need not be activated automatically (if visual attention is directed elsewhere, say) or involved at all (if the object is not even within view). In general, perception may operate slowly, bringing news of the object's position that could have been *predicted* from the beginning of the motoric command. Such prediction may be of independent value, even if one still *could* perceive the position after the fact). Active perception (which bears a cost upon consciousness) may then be used moderately, or operate more smoothly, aided top down by what to expect and only adjusting in case of miss-prediction. But all of these instances would work only given that the position the object is moved to (as an informational item managed by the manipulation procedure) is regularly where it would be perceived to be (the result of the query procedure).

We cannot exclusively keep just to collections of procedures that are object-coherent. Sometimes the environmental or cognitive regularities may change, which then breaks down previously valid arrays. Some constructions (as we see next), even taking in just valid procedure arrays, may produce object-incoherent ones. I thus introduce the following concept:

Coherentation: an operation that, given an object-incoherent collection of procedures, produces a procedure array that is the “natural” or “nearest” object-coherent resolution.

I leave the notion imprecise, but I use it only in cases where the meaning is hopefully clear. For our simple illustrative cases, this should suffice. Note the following:

---

[53] For comparison with robotics, (Modayil, et al., 2008), too, is "inspired by theories in child development, in particular the assumption that coherent motion is one of the primary mechanisms for the initial perception of objects". Motion is a central form of physical change, rooted in the spatial. But from our perspective, generalized beyond the spatial, the deeper reason underlying its primacy is coherence writ large, abstractly, not just spatiotemporally. (This is also related to the issue of prediction, highly debated as a unifying fundamental concept in cognitive science. But I do not address that.)

i. A reasonable coherentation need not always exist. I keep the usage to resolutions that are near enough; that keep to the intent or core of the incoherent collection.
ii. There may be more than one reasonable resolution.
iii. As object-coherence is a very strong, semantic form of consistency, so is coherentation: It is dependent on the environment in which the system is embedded, on successfully capturing all relevant regularities and finding a resolution within them.

An array, then, *determines* the *facts* about it, about the interaction between its procedures, in relation to the objects. For example, it is a fact of apples as an object-type, that we cannot eat the same apple twice. The facts, however, are not *part* of the array; they are by no means (cognitively) *in* it. The *apples* object-type may assimilate the independent facts that 1) eating an apple leaves just its core (for example by expecting this result of its perception), and that 2) an apple's core is inedible, but not the implications of the composition of the two. At the very highest levels, some meta-procedure may not only take into account, but perhaps even "perceive" such facts in some sense. But it is atypical for such a meta-procedure to be part of the array. When we reflect upon the object this way, we do something very different from *using* it as an object. And mostly, understanding most facts (for sufficiently rich arrays) may simply be beyond our capabilities. The facts, not themselves represented in the system, certainly do not *constitute* the object-types (as with logicism, Section 1.1, and Model Theory's types, Section 6.1.1); they may not necessarily even determine the types. Part of these facts are what properties can be coherently *attributed* to the object-type, which is constituted upon the array. I do not require that these properties must actually be implemented by the array, explicitly. For example, the procedure array for handling rectangular boxes may perceive and maintain a box's three edge lengths without having any computational element for maintaining its volume or a procedure for calculating it. Its volume may be *deducible* from this representation without itself being represented. Operations can be *as if* they are about the non-represented property (e.g., using the box to store some fluid from a differently shaped vessel). The more central such a property is, the more efficient it may be to actually represent it, architecturally. But the richness of the object-type is not restricted by the aspects of it that are *actually* represented; just by those coherently determined by its procedures.

Let us leave aside the deep metaphysics lurking here. Importantly, we have no need to explore the details of *how* a procedure array comes to give rise to its type. For example, whether and how two operationally different arrays, as real-world processes, may bring about the same type[54], an abstract notion, and whether the system itself can notice this difference or whether it can relate in any way to the abstract entity rather than just its concrete representatives.

6.3 Extensions & Amalgamations

For understanding object *constitution*, as a cognitive activity, we cannot make do with a snapshot of this cognitive framework (the arrays and the object-types they determine and operate on). We need to address its developmental track: how the system comes to attribute objecthood and also alter it, enriching object-types when it learns that it can and breaking down others when it learns that it should. A full account would require attending to continuous changes in (the parameters of) procedures of an array – under its coherence constraint, to possible abrupt switches between alternative, incompatible coherentations and more. But for my particular mathematical application, it suffices to account for how procedure arrays extend others (as partially ordered sets), with respect to the object-types associated with them. For that, I introduce the concepts of *extension* and *amalgamation*:

Extension: An extension of a procedure array P_A is a set of procedures $P_B \supset P_A$ that is itself object-coherent; i.e., an array. (P_B is then said to *extend* P_A).[55] The terminology transfers to the object-types brought about by the arrays too: T_B extends or is an extension of T_A, respectively. (And in the opposite direction, P_A and T_A are *reducts* of P_B and T_B, respectively).

An extension is a less abstract object-type that allows for doing more. Hence, an extension might allow for fewer objects to be perceived *as it* (i.e., it determines a subclass or sub-category). An infant might have a *generic* object-type of $fruit_g$, that she has come to expect to find on her plate at dinner for dessert. Later on, she may learn to distinguish the particular type that is *orange*, and master not only the expectation for its

[54] This is related to – but not the same as – the notion of *bi-simulation*.
[55] The familiar analog in Object Oriented Programming is inheritance.

particular taste – but the art of peeling it. *Orange*, having these extra procedures, extends that general $fruit_g$. Later on, when she has mastered many common types of fruit, a *unified* object-type of $fruit_u$ may thus be cognitively formed, which incorporates these mastered object-types. Importantly, this incorporation is now *not* an extension; it is not, strictly defined, set containment. To be sure, the child can now independently eat a piece of fruit, even if it happens to be an orange, by peeling it first as needed. However, "if the fruit is an orange, peel it"[56] is a different procedure, which uses (calls as a function), builds on the previously-mastered "peel the orange" one. This meta-procedure, "Peel the object", is simply not one of the procedures that applies to all fruit[57]. So, while the original, generic $fruit_g$ may be thought of as the *intersection* of all specific fruit types (their "common denominator"), the latter $fruit_u$ is not formally their *union*. This latter notion of "extending" object-types (from oranges to $fruit_u$) by *generalizing* them, extending the domain without abstracting, is not required herein and thus is not pursued further. On the other hand, the gap between $fruit_g$ and $fruit_u$ does exemplify a critical divergence of (my rendition of) *cognitive* object-types from standard worldly types, that seem to not depend on our conceptualization of them and its development but stay fixed (as just the category of *fruit*).

To make explicit the main relation between an extended type and the type it extends: For $P_A \subset P_B$ as above, the extended object-type T_B can in particular restrictedly be thought of, and cognitively acted upon, in terms of T_A: an operation $O_A \in P_A \subset P_B$ on T_B, valid for T_A (i.e., does not break its coherence), is in the extended P_B array and thus by definition still valid. [58] A construct of P_A (e.g., box volume, even if it was not explicitly represented) could be constructed of P_B just as well. A fact about a specific procedure in P_A or the interactions among some of them is still true of them as embedded in P_B (though not necessarily for facts that also involve P_B, such as the type as a whole).

---

[56] or more likely, the general "if the fruit is an orange, perceive it as an orange instead of a $fruit_u$, and continue". (But I do not explain how the system perceives an object as one object-type rather than another).

[57] Even if it technically were (i.e., without strawberries, etc. in the world), then even though technically a watermelon (say) could be peeled, it would not be relevant as a cognitive object-type, for which mastering the preparation of a watermelon typically goes through other means.

[58] The following require some care when taking a coherentation, which may mess things up. E.g. some constructs become ill-defined.

Adding a particular procedure (or a few) to an array, learning a new thing that one can do to a type of object (bite an apple) or notice about it (the apple is rotten!), is a fundamental learning operation of the objectification system. However, there may be importance and value also in bringing together larger collections of procedures – when the added collection is not of just an arbitrary selection of procedures (that may likely not produce object-coherence and be hard to coherentize), but a previously designated, internally coherent one. This brings us to an (also rather fundamental) generalization of extensions, which reflects the complexities of extensions too, and is central to my intended applications. It aims to capture the case of object-types that are built on a *combination* of more than one simpler object-type: Bats, which are both mammals and flying animals (though not necessarily the only ones); a commander, who is a "combination" of a boss and a military person; a visually perceived object, which may perhaps be amalgamated from each eye's own perceived object (whose location in 3D is less determined); various other visual integrations and also *multi*modal integrations, perhaps; a smartphone (Fig. 4)[59] ; and much more. Formally, then:

Fig. 4 The smartphone as an amalgamation
Source: The Internet

Amalgamation[60]: Given two (or more) procedure arrays P_A, P_B, an ***amalgamation*** of them is a set of procedures $P_C \supset P_A \cup P_B$ that is itself object-coherent, i.e. an array. The terminology, again, transfers to the object-types brought about by the arrays too, such that T_C is an amalgamation of T_A and T_B. [61]

[59] Thanks to Neta Livneh for this example.
[60] In model theory, the *amalgamation* property concerns (roughly speaking) two structures that may be regarded as substructures of a larger one.
[61] Compare with *conceptual blend*, "the conceptual combination of two distinct cognitive structures…" (Lakoff, et al., 2000), section 3.3.2.4.

Importantly, the set union itself might ordinarily *not* be object-coherent (as we shall soon see).

Both amalgamated arrays P_A and P_B are extended by the amalgamation P_C, and so the same reasoning as above is applicable: T_C can in particular be restrictedly thought of, and cognitively acted upon, in terms of T_A or T_B: an operation from P_A or P_B belongs to P_C too (and does not break its object-coherence). [62] A construct of P_A or of P_B could be constructed of P_C just as well. And finally, a fact regarding a specific procedure in P_A or in P_B, or the interactions between procedures from P_A or between procedures from P_B, is still true of them as embedded in P_C (mixing procedures from P_A and P_B, however, can change things). [63]

There is much more to (extensions and) amalgamations than that. Let us now explore the concept at some depth, with regard to arrays that each concern a *particular* object. Much of the complexity (relevant for us) can already be found there. For starters, then, consider procedure arrays that are totally independent, with no interaction between them (cognitive or physical), concerning, as it were, different "dimensions". Their union is a trivial amalgamation that can be thought of as their *Cartesian product* (which, as with its abstract category-theory formulation, need not exist). For the sake of illustration, let us take dimensions literally: Consider a bead that can be moved freely on the unit, $[0,1]^2$ square board, and two arrays P_X and P_Y that allow for moving and perceiving the bead's position at their respective dimension (Fig. 5). Part of each array's object-coherence is that moving some distance in some direction is correlated with its perception as having moved accordingly. The arrays' union is an amalgamation, P_{XY}. It determines a unified type, now with two data points, that are simply kept separate; with no interaction between the respective dimensions. For that matter, it could all be about two different beads, on two separate wires, one vertical and one horizontal, each

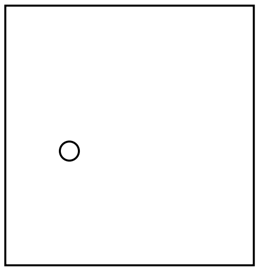

Fig. 5 Amalgamation of two non-interacting arrays (for object movement)

[62] Like for extensions in general, coherentations may convolute the following.
[63] The analog in OOP is known as multiple inheritance.

independently perceived and manipulated. The array is not rich enough to distinguish between these two cases.

Where this notion starts to get interesting[64] is when there are *interactions* between the arrays (which for simplicity I assume are implemented independently in the cognitive system). When there is *regularity* to the interaction, this is where amalgamations get their substance. By assimilating such regularities is how they earn their keep; usually, with added procedures that are not in either of the original arrays, but that relate *between* the originals' procedures (and between their constructs). When properly conforming to the regularities, the added procedures can still come to produce a proper, object-coherent array. Such an amalgamation determines a *higher-level* object-type. I now exemplify and elaborate.

Consider the example above of a bead moved freely on a board, where now it is instead on a wire, constrained to moving only along the x = y diagonal of the same unit board (Fig. 6). The procedures of P_X and P_Y are the same, but now in reality, moving the bead along the x-axis also changes its y-position (and its perception) and vice versa. The Cartesian product of the arrays, P_{XY}, fails to reflect this regularity. Subsequently, the object-type it could have determined would be too general, allowing for situations that never happen ($x \neq y$). More importantly, movement along the x-axis would no longer be properly coordinated with perception along the y-axis (and vice versa) – which would break its object-coherence.

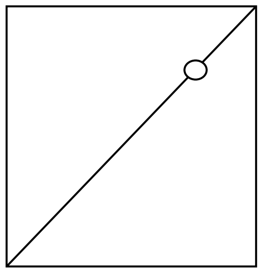

Fig. 6 Amalgamation of two arrays (for object movement) with interaction x=y

Enter $P_{XY}D$ instead, the coherentation of the product relative to this added background regularity (that is the Diagonal). Where P_{XY} had two distinct position components, $P_{XY}D$ now has an implementation that links the two. For example, say an x-movement (in P_X) changes the ("expected") y-position (perception must not be activated automatically, or it could function more slowly, only fixing in case of miss-prediction). Then in $P_{XY}D$ it would come to change the y-position too, accordingly.[65]

[64] It is this point at which we break off from OOP (in my understanding). In OOP, multiple inherited classes cannot interact but must be independent to begin with. Only then can further connections between the two be coded in. But again, for us, interactions can preexist – and must be learned, fundamentally.

[65] "A conceptual blend is the conceptual combination of two distinct cognitive structures *with fixed correspondences between them*" (my emphasis); (Lakoff, et al., 2000), section 3.3.2.4.

The two position components, now perfectly linked, effectively function as a single parameter, and thus could actually be unified into a single one, become a more efficient conceptual and computational entity. $P_{XY}D$, though starting out not only with double the procedures but double the data elements of P_X or P_Y, does not become more complicated in terms of the possible situations ("ontology") it needs to represent and manage.

Compression is about much more than just saving up on some neurons. To achieve it is to break down how the world is and how to exist in it. New possibilities for how to operate arise: the y-position can be manipulated through both manipulation procedures, and likewise for the x-position. More generally, this stands for the query procedures (of position-perception), and most generally, for *any* combination of procedures from $P_{XY}D$. It is possible to y-move the bead down by 0.3, x-perceive its position, and, according to whether it is left or right of the middle, y-move it left or right by 0.2.

This very simple amalgamation captures something fundamental about the objectification of physical objects, located in physical space-time. The visual perception of an object, its auditory perception, sending a hand towards it, etc. each have their own spatiotemporal location. However, these must all be coordinated, not essentially differently from how the visual perception of on object's particular parts need be. Each one's own dimensions, in extension or amalgamation, should all collapse into the same ones, along the lines of what is expressed in Fig. 6. This is a fundamental space-time *manifestation* of object-coherence, that for physical objects, the system needs to figure out.

For the previous simple diagonal case, the ability to act upon the object freely using procedures from both P_X and P_Y does not actually contribute anything; everything is just reflected in the other dimension. $P_{XY}D$'s object-type, $T_{XY}D$, *de facto* now has only one component, of position along both axes. (And indeed, the system may optimize the implementation accordingly). Where things get interesting, is when new feats can be achieved; through a combination of operations, new *affordances*[66] for the object-type

[66] Gibson's original definition was for objective possibilities, within the environment itself. But given our focus here, on cognitive object-types, I use it in Norman's sense, as *perceived* possibilities. (Wiki-Afford, 2018)

arise. Consider a somewhat richer setting, where instead of on a diagonal, the bead moves along a Quarter-Circle wire between (0,0) and (1,1). The product coherentation, $P_{XY}QC$, operates quite similarly. At the level of implementation, although a single data component could still suffice, it may now make sense to keep to two that are synchronized. What is important here is that the interaction between the two arrays' procedures must no longer be trivial as before. Say, for example, that there is only a four-level (uniform) sensitivity for each axis' precision (Fig. 7). Then, $P_{XY}QC$ now allows for a better perception than each of P_X or P_Y, using x-perception for the lower 45° and y-perception for the higher. The added richness comes at a price. The less-trivial regularity (between the dimensions, and so between the arrays) is harder to pick up, and with it, on the potential for amalgamating. The coherentation is more drastic, which renders the computational management of the interaction between the two arrays difficult. These difficulties, however, are essential to an important task that evolution takes on – 3D vision:

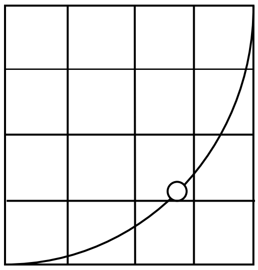

Fig. 7 Amalgamation of two arrays for object movement in four-partition with interaction along the quarter-circle.

The basic setting is quite similar to this last example: Two eyes share a field of vision and watch the same scene (and to stress the similarity, let us ignore their vertical perception). But now, they do so from an almost identical angle. Unlike in our first, Cartesian product example (Fig. 5), neither of them alone can perceive an object's *depth*. But their amalgamation can *use* this delicate and intricate interaction between their horizontal perceptions to do just that. The usual exchange applies: - "Being so fundamental, it *could* all be innately coded in"; - "In principle, yes, but as long as we have such a system for other purposes anyway, why not just use it – and moreover gain the advantages in coding just general principles for adaptation" (Section 6.1.3). This very last point can be made into a (sort of an) argument: Binocular vision is a more difficult task than just two-eyed vision in which each eye is directed sideways at its own field of vision. It thus have probably developed later on in evolution. How could that come about? The reasonable path[67] for this evolutionary development would be that a mutation brought the two eyes, and hence their fields of vision, closer together

[67] Suggested by Aaron Sloman (private communication)

(against the selection for covering together as large a field of vision as possible). This conjoining would only confuse an *innate* system that has yet to adapt to binocular vision. But for an open system, able to relate such interactions as they arise, this would, just the opposite – be selected for. The alternatives would be that the cognitive system adapted to binocular vision *before* or *with* the eyes actually gaining a shared field of vision, which are both highly unlikely possibilities. To sum up this (repeated) point: objectification in the abstract is most general and thus relevant to many aspects of cognition, high and low – and thus likely to be implemented in the cognitive system, with that objectification system used broadly.

To stress the centrality of amalgamations for creating higher-level object-types (and in the following case objects) that offer new affordances, I offer another example of a different nature: the smartphone (Fig. 4). It incorporates several types of devices (that also are, or used to be, separate ones); it can be used as each of them. But they can also interact (through the single computer at its core). So now, even though the GPS tracker was put in particularly for the intended purpose of making the device a navigation system too, it can interact with the camera sensor (which was put there for a different, also-particular purpose), to produce pictures that incorporate the location at which they were taken (two very different data types). A general point is the following: amalgamating *different* but interacting modalities or "dimensions", may be harder to achieve (to coherentize) – but proportionally to the value that it affords. We now return to the complexities that can be illustrated even through our series of simple examples.

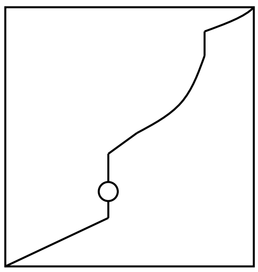

Fig. 8 Amalgamation of two arrays for object movement with non-1-1 interaction (multiple y's per x).

Considerations similar to the bead on a quarter-circle wire above (Fig. 7) would generalize for monotonically rising wires in general. Now consider instead a wire that at some sections simply climbs vertically (for various fixed x-values) (Fig. 8). The wire still serves as a one-dimensional constraint upon the bead. But now its full x-y position cannot be "thought of" (perceived or manipulated) just in terms of its x-position. Whether y-moving it would change its x-location may depend on its y-position, an aspect to which the x-perception is blind (if the bead happens to

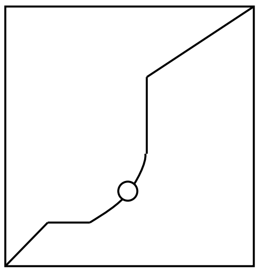

Fig. 9 Amalgamation of two arrays for object movement with general non-1-1 interaction

be on the vertical section), and so these are no longer fully coordinated. We cannot switch from the object as a y-object to it as an x-object and then back again, as the latter may be ill defined. It can still be managed as a y-object, but that too would collapse in case of a wire with perfectly horizontal sections too (Fig. 9).

To be sure, a richer amalgamation could fit in both original arrays and manage the (wire) regularity between them. But in order to do so, it needs to maintain (and continue managing) both of the original location parameters. More generally, it also must keep the original arrays separate. To illustrate through a somewhat different example: Consider two arrays, one (P_W) for recognizing words (i.e., supporting reading), and another (P_{DS}) for recognizing decoration styles, regardless of the written semantic content. Consider also a toddler (who has mastered P_{DS} but has yet to master P_W) looking at some inscription of a word (W_1) in a fancy font. The inscription, not making sense, is perceived as a bearer of decoration and no more. In particular, were some other word (W_2) written there instead, it would not necessarily be differentiated. (In full actuality, a change happening right before her eyes would be noticed, of course. This would be through the far richer array supporting visual perceptions of things in general, which are neither writings nor decoration styles. Even so, upon perceiving W_1 and then *later on* seeing W_2, separately, the change may certainly go unnoticed). There is an intermediate mastering stage, where the preschooler already understands that different inscriptions can have different meanings – somehow. But consider the grown-up child who has completely mastered P_W too (and reading in general) and can now tell apart *different words written in the same font*. She can keep to the original arrays, and also amalgamate them to treat the text and the decoration style as two aspects of inscriptions as a richer object-type. What she can no longer do is hold on to inscriptions as a unified object-type for which the decoration-style is *identified* with informational content, alone determining the objects that may be (ontology).

Finally, let us consider another related commonplace example for which the original arrays differ (and do not collapse onto the same dimension). The setting is an object in front of a woman and two related procedure arrays. For queries, both arrays let the woman perceive the object's angle towards her. For manipulations, the first lets her turn the object (left or right) with her hands, while the other lets her move around it (always facing it). Their Cartesian product amalgamation would simply duplicate the perception and contain both manipulations while keeping them separate. That is, there would be two empirically but not cognitively identical angle parameters. But the first array's perception does not relate to the second's manipulation (of moving around the object) and vice versa – and so object-coherence would be lost for the combined collection of procedures.

Now enter the coherentation of that amalgamation (relative to the relevant real-world geometric regularities). It lets her do both manipulations, but also "assimilates" the relation between them: She can now turn the object to the right, turn around it by the same amount, and have the perception of its identical angle towards her match her expectation. The amalgamation's type is far more enriched: For each of the original arrays, keeping the angle between her and the object had only the trivial solution. But for this amalgamation, there is a whole circling dimension of solutions. For each array, there was the construct of "the angle between me and the object", which was the same as either "the angle between the object and the world" or "the angle between me and the world" respectively. Now, though, the construct of "the orientation in the world of the angle between me and the object" (though non-intuitively sophisticated) is no longer trivial (its richness, part of the coordination between the manipulation and query procedures).

Some important words of caution regarding coherentations are required:

1. This last example as well as previous ones may seem to suggest that for a minimal amalgamation (of which others are an extension), we (conceptually) and the system (operationally) should always just take *coherentation* of the procedures' union. However: Coherentation is dependent on the regularities, which are not inherent to the original arrays themselves (and the regularities that support their own object-coherence). This means that what "the minimal amalgamation" is can change according to the context – its scope and its dynamics. This phenomenon stands at the heart of Chapter 7.
2. The level of informality at which I have left the concept spreads further, infiltrating our seemingly more formal conceptualizations: Altering the procedures means that a *coherentized* amalgamation (or extension) is no longer formally an extension of the original arrays. A further account would be required, e.g. for specifying when an altered procedure was not *essentially* altered.

3. With *coherentation* built on *natural* or *nearest* (object-coherent array), the result of amalgamations, extensions, and coherentations, may well depend on the way we break it down as a series of consecutive operations, and on its particular order.
4. Some types of objects, i.e. categories (of the objects in the real world), are nicely ordered in a hierarchical fashion. *Winged mammals* are all of the objects that are both *mammals* and *winged*. If all there was to the cognitive space of object-types was the partial ordering of sets of arrays, this would suggest that all relations between categories are like that. It would suggest a nice, orderly matched ontological hierarchy in the world (as conceived), from cognitive extensions and amalgamations to subsets of objects. We know, though from Rosch's seminal works and on (Section 6.1.4), that many natural or cognitive categories are nothing like that; and that this is in fact the general rule. This, I suggest, is precisely because the space of objectification itself, led by "distorting" coherentations that bring in the world and its rich interactions, is itself not really ordered that way. Extensions and amalgamations are still important as basic relations and constructions, i.e. ontologically and developmentally. But the richness of reality penetrates the absorbed interactions.

For my limited purposes, there is no need to unravel these complications. So, pragmatically, instead, I leave them as warnings to bear in mind (and as central issues in taking this approach to objectification further).

There is another general theme that arises from this last example and the previous ones. The free bead on the table had more positions to occupy (or to reframe for future purposes, "the *positions* as a type had a more numerous ontology"). But the one on the quarter-circle, though restricted to one dimension, was of a higher-level type, with a non-trivial interaction between its components. The general rule is the following: The more procedures an array has – the richer its object-type is, with more available perceptions and manipulations to it. However, the more interrelations there are *between* its procedures and the more regular these are – the higher-level the object-type. (Compare Fig. 5 < Fig. 9 < Fig. 8 < Fig. 6 ≈ Fig. 7). An amalgamation always produces a richer object-type. But a more substantial coherentation, necessitated by interrelations between its original arrays – produces a higher level object-type.

For the sake of illustration, most of these examples used two rather identical arrays. But the demonstrated principles are (claimed to be) completely general. Moreover, they pertain more simply to extensions, too, in which there is only one preexisting type of entity (but I skip a rehash of the details – of how adding an interacting procedure to an array may break coherence and calls for a coherentation, etc.). Much more generally, the principles pertain to the much more common type of object-types – that support a *multitude of objects*. But accounting for this requires diving into further issues that are fundamental to objectification: the detection of objects *as* a particular and valid object-type ("- Is that an apple (/ fruit)? - No, it's your tennis ball!"); tracking them ("That dog's running off with your tennis ball!"); and assimilating and representing the interactions between *them* ("Forget it, it was a silly bet anyway, you were never going to stack it on top of that other ball"). I merely assume that all of these can be captured properly in an extended framework.

6.4 Regularities

Regularities are what enable a reduction of the data space (the results of queries), and an increase in control over it through control of whatever actual space of which the data represents information (the results of manipulations). Within an a priori setting of queries and manipulations, regularities are what then structures ontology. It is up to the system to grasp and appropriately coherentize them. As usual, this includes not just natural cognitive systems, but also artificial ones. For example:

> The robot learns actions by observing the effects of performing random motor babbling in the presence of the object. Motor babbling is a process of repeatedly performing a random motor command for a short duration. One strength of the action learning approach we present here is that the robot is able to use this goal-free experience to form actions that can be later used for goal directed planning. (Modayil, et al., 2008)

(Quite substantial analogies to human development can be found in developmental robotics. See for example (Ugur, et al., 2011)).

When the system has no clue, all it can do is act randomly (spanning the whole space of possibilities), observe, and try to learn. Such "babbling" is an optimization that facilitates the detection of the regularities between the procedures (both perceptions and actions).

Regularities, however, can be partial. How about when throwing a ball high up in the air isn't quite coordinated with catching it? (For an array that supports only these but not moving, nor perceiving it outside the vertical axis). What happens to object constancy? The framework I offer is a limited one, focused only on approximately perfect regularities, which can be treated non-probabilistically. Being concerned with mathematical objects, this is all that is required for its application. Of course, it is all embedded within the much larger picture with which the cognitive system is concerned; and it is this general context in which the system is endowed with the capacity to detect approximately perfect regularities too.

Even approximately perfect correlations come in many forms. A regularity may be a mathematical necessity, a physical one, biological, a contingency of evolutionary history, a cultural construct, etc.… or just an environmental contingency, the way empirical statistics happened to be (or are biased to be, for example, by teachers). The system, in general, cannot, nor for the most part need not, distinguish among of these options. It is unable to tell apart contingently perfect correlation from causation and determine the cause, its direction, and its depth. And that is fine, it only has to pick up on the regularity, not do metaphysics (the value of which is so much more indirect).

To demonstrate how even just environmental contingency can determine coherence, consider, first, the visual system's heuristic usage of shadows for figuring out 3D location of objects (Fig. 10). The perceptions of the ball and its shadow are coordinated

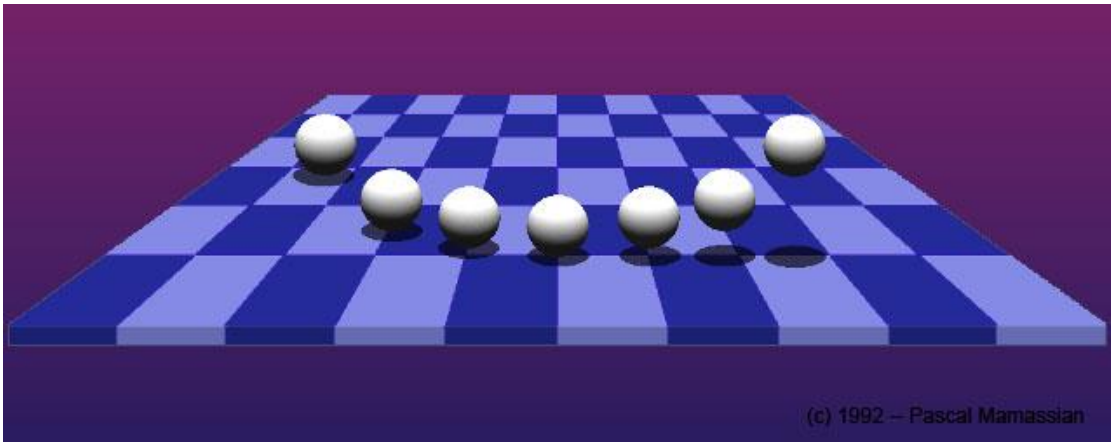

Fig. 10 Shadows tending to appear below objects is an environmental, physically non-necessary regularity. Visual perception uses this non-consciously detected regularity to tell apart depth from height in 3D space.

and together support a higher-level construct (than each one's vertical position in 2D), which arbitrates in 3D between the ball's height and depth (a construct which could then be coordinated with other procedures, such as grabbing the ball). In this extended sense, the ball's shadow is part of the ball, just as a part of the ball is a part of the ball (if with various finer differences, e.g., the shadow's non-necessity). Described the other way around, the amalgamation that relates (the perceptions of) the shade's position to the ball's, could be made coherent only in the face of this regularity. Without it, only the lower-level amalgamated object-types, kept separate (and thus not allowing for positioning the ball along the depth-height diagonal axis), are possible. Perhaps this shade-heuristic in particular happens to be innate (involving the most basic issues for a living being). Or perhaps we are flexible enough to adjust to a world with multiple light sources (say). For a more clear-cut case about this, let us (for a second example) reinterpret a historical classic:

> Pavlov became interested in studying reflexes when he saw that the dogs drooled without the proper stimulus. Although no food was in sight, their saliva still dribbled. It turned out that the dogs were reacting to lab coats. Every time the dogs were served food, the person who served the food was wearing a lab coat. Therefore, the dogs reacted as if food was on its way whenever they saw a lab coat.[68]

From there on, Pavlov famously made his dogs associate the sounds of a bell with food. And so begat *classical conditioning*[69], a pillar of behaviorism (alongside Skinner's *operant conditioning*; Section 1.1). From our cognitive, representational point of view, we can say that food, as an object-type, came to be perceived not just by its (naturally occurring) looks or smells, but through other stimuli that (should have) had no association with it. The objectification system simply detected these regularities (whatever their origins or cause), and constructed the object-type accordingly. And if the system can do this even for the most fundamental objects in life (as food is), with the leading candidates for being innate rather than acquired, then the scope of

---

[68] "Pavlov's Dog". Nobelprize.org. Nobel Media AB 2014.
http://www.nobelprize.org/educational/medicine/pavlov/readmore.html

[69] The pairing of a conditioned stimulus with an unconditioned stimulus.

objectification must be considerable indeed.[70] If it can do that even across modalities, for example, then issues explored in cross-modal integration cannot be ruled out a priori from being managed by the objectification system, perhaps even as amalgamations rather than just a gradual extension from primary modalities to other ones. At this level of abstraction, they certainly can be made to fit this conceptualization:

> Through detailed long-term study of the neurophysiology of the superior colliculus, [a "groundbreaking work in the modern field of multisensory integration"] distilled three general principles by which multisensory integration may best be described.
>
> - **The spatial rule**[71,72] states that multisensory integration is more likely or stronger when the constituent unisensory stimuli arise from approximately the same location.
> - **The temporal rule**[72,73] states that multisensory integration is more likely or stronger when the constituent unisensory stimuli arise at approximately the same time. (Wiki-MSInt, 2018)
>
> [The third rule is for the next section].

The space (in the abstract sense) in which *physical* objects are embedded is physical space-time. The interactions between various (perception and action) operations on objects of this kind – pass through their spatiotemporal *locations*. This makes the spatiotemporal location the nexus of the *regularities* to those interactions. This is where the system could pick up on the structure of the interactions (that become more complex, more irregular, with distance). This is where, furthermore, the system would be better off innately guided (determined or biased) to look for regularities (but the latter is an optimization rather than a requirement).

Such empirical findings can thus be taken for concrete *manifestations* of the abstract notions, of object-coherence, coherentations, and regularity detections. Other concrete

---

[70] Though as for behaviorism, such associations in and of themselves do not necessitate any such cognitive representations to begin with.

[71] Meredith, MA.; Stein, BE. (Feb 1986). "Spatial factors determine the activity of multisensory neurons in cat superior colliculus". Brain Res. **365**(2): 350–4.

[72] King AJ, Palmer AR (1985). "Integration of visual and auditory information in bimodal neurones in the guinea-pig superior colliculus". Exp. Brain. Res. **60** (3): 492–500.

[73] Meredith, MA.; Nemitz, JW.; Stein, BE. (Oct 1987). "Determinants of multisensory integration in superior colliculus neurons. I. Temporal factors". J. Neurosci. **7** (10): 3215–29.

manifestations, at many different levels, may be found or could be looked for or conceptualized as such, as manifestations of these principles. At higher levels (than, for instance, multi-sensory integration, which is biologically fundamental), guiding the system on where to look for regularities shifts from the innate to more adaptable means. *Language* is a particularly important tool here. Generally, *linguistic distinction* can aid in noticing and assimilating finer distinctions within the perceived ontology. In the opposing direction, *linguistic designation* can help bring things together. This involves not only various biological modalities (which perhaps suggests to the infant that the important object to focus on representationally is *mom*, not separately *mom's voice* and *mom's face*). It can involve highly conceptual cross-"dimensional" interactions, even social ones. *School*, for the youngster, is not at all just a physical area ("Are you at your school?"). It comes to be integrated with various people, activities, norms, etc. ("Are you at school", "How was school", etc.). Part of this is taught explicitly ("This is your school", "you have to go to school"); but much of it can be the rich, if less-defining, regularities that come to be detected and coordinated through linguistic designation ("How was school" as reference to both classes and recess). Language, predetermined by past detections (and creations) of object-types, guides the objectification system in choosing procedure sets, detecting the regularities, and constructing the object-types.

Without, or at least before, delving into the intricacies of innate versus learned knowledge, what matters here, abstractly, is these general principles of objectification, and that they actually *are* handled – somehow – by the cognitive system, given that we are able to learn new object-types and come to master them. That this includes learnable mathematical ontologies is my general assumption.

6.5 Development

I have presented a basic formal specification for cognitive objectification. It adopted a (suggestively) incremental view, in which (already coherent) object-types can be combined appropriately into richer ones. Though very partial, the framework intends to capture an aspect of how the cognitive system actually functions. An actual system needs to support, bring about, and manage something like what I have abstractly specified: The arrays, their extensions, amalgamations, and coherentations (extended for multi-object object-types). For this, in particular, it needs to recognize the regularities between procedures (on its way to approximately reconstructing real-world objecthood). With all of these aspect in place, we have a basic grip on the *structure of the space of possible objectifications*. But for the broader view, we need to understand – and the system needs to implement – the developmental dynamics: how it moves about as a set of points in this space, through arrays that come to be actualized and used.

The full space is structured as a landscape, where some sets of procedures are of higher *value* than others. Our first, basic factor was object-coherence, without which the ability to use that object-representation for prediction and manipulation of the objects collapses. The arrays are thus local attractors (driven to through coherentations). Everything from here on is about what the best sets of peaks are and how to find them. What other factors are there?

Before another step, let us rule out the trivial solution. Why not just take *all* possible arrays as object-types of the system? Is objectification as a learning process simply the quest to find, within the space of procedure-groupings, the coherent ones? In order to serve its purpose, the system must be *economical* (a general leitmotif for this section). It must keep low the number of object-types it maintains and manages. In the end, at a particular instant, we perceive something mainly as a *particular* object-type. This is what frames our perception, focusing our attention on the important properties and affordances for further predictions and manipulations. (E.g., we need to perceive an orange as an orange rather than a fruit in order to realize *how* to make it edible; but we need to perceive it as a fruit instead in order to know that it would satisfy what we were asked to bring). Too many options for the object-type would defeat the purpose, of finding the *most useful* object-type and representing the object through *it*. (Switching between object-types may be costly, and selecting the right one is another burden on the system even if always done right). To some extent, we can influence (i.e. top-down) the choice of object-type (out of a few possible ones); but that, too, would be defeated by an overload of possibilities. This is thus a proper optimization problem. There must be, and there indeed is, much more to the value of an object-type than only its coherence. Let us now map some central considerations.

The central issue we've been exploring is that of extensions. How does it relate to the value of a type, then?

To be sure: There is nothing to be gained by bringing together many unrelated procedures just for the sake of it, extending and amalgamating them for merely *richer* object-types. (Cartesian products were brought in as a formality, being behaviorally isomorphic to the unamalgamated arrays, not contributing anything to the system's abilities on its own, and that need not be found in the actual system). In fact, richer types are a computational burden on the system, which is forced to attend simultaneously to different aspects of the richer object and what it allows for, where perhaps a meager type would have sufficed.

However: There is much to be gained from having together many *substantially-interacting* procedures, for the sake of forming *high-level* object-types (section 6.3). This is what allows for the integration of perceived inter-procedure regularities into the ontology, and thus enables substantial predictions and reveals affordances. This is how the system tracks the structure of reality while keeping its representation dimensionality in check (a compression that reflects comprehension). Being *driven towards high-level object-types* (higher peaks) is the fundamental (and hopefully reasonable, if not obvious) principle that I offer, as the underlying form that the system's progression as it acquires new object-types takes.

The drive for high-level object-types is *as a by-product* also a drive for *rich* arrays, which thus needs to be weighed against its added cost (a comparison that may be difficult to characterize in the abstract).[74] Moreover, high-order-ness of an object-type itself, alongside its representational value, also has its own computational price (involving more procedures with more interactions). Whether a higher-order object-type is worth these two costs, its value, depends on the actual degree (recall that Fig. 9 < Fig. 7); on how regular the interaction is, thus giving combinational value and compressibility. Weak connections might sometimes not be worth the effort of

[74] The physiological procedures (photoreceptors etc.) are here just accepted as given. But on an evolutionary scale, they are certainly manageable, predetermining the individual organism's innately fixed atomic procedures – according to similar considerations. For example, the development of eyesight in complexity proportional to its value for a creature of that species (including actually deteriorating eyesight when finding a niche underground, rendering it unuseful). As for non-peripheral, meta-procedures, however, the individual may well have some flexibility in forming new ones, rather than just learning object-types on the basis of a predetermined set.

detecting and managing them. This is reflected empirically for example (continuing p. 162) in the third of the general principles of multisensory integration:

- **The principle of inverse effectiveness**[75,76] states that multisensory integration is more likely or stronger when the constituent unisensory stimuli evoke relatively weak responses when presented in isolation. (Wiki-MSInt, 2018)

The implications of the *cost* of the high level are much more general than that. Approximating reality as fully as possible, though fundamental, is not everything. The abstract matters too (and to belittle its value in a dissertation on mathematics would be preposterous). Consider a classic of the cognitive science, with regard to the value of different object-levels and the selection between them:

> Rosch et al. (1976) had argued that category knowledge is organized hierarchically into various levels and the three that are most often discussed are the superordinate, basic, and subordinate levels. For example, subordinate categories like beagle and boxer are nested under the basic-level category of dog, which in turn lies below the superordinate categories of mammal and animal. Rosch et al. also claimed that the basic level in the hierarchy is the most functional: the level at which we do most of our thinking. Basic-level superiority occurs because members of basic-level categories tend to possess significant numbers of attributes in common, have similar shapes, and invoke similar motor movements. (Goswami, 2010 p. 133)

The explanation here for a level's superiority is in line with my own framing in terms of the value of the high-level (e.g. *orange* over $fruit_g$). But this only explains why *dog* is (usually) superior to the "superordinate category" (to use their terminology here) that is *animal*. In fact, the very same reasoning (of tending to possess significant numbers of attributes in common etc.) would actually further still favor subordinate categories like *beagle* and *boxer* over *dog*. Yet, these are *not* the so-called *basic* categories that we select and use by default. The "most functional" level, the one "at which we do most of our thinking", simply is not always (and ordinarily is not) the one that offers *the most*

---

[75] Meredith MA, Stein BE (July 1983). "Interactions among converging sensory inputs in the superior colliculus". Science. **221** (4608): 389–91.
[76] *Meredith, MA.; Stein, BE. (Sep 1986). "Visual, auditory, and somatosensory convergence on cells in superior colliculus results in multisensory integration". J Neurophysiol.* ***56*** *(3): 640–62.*

function and facilitates the most thinking (with regard to the object). Let us dwell on that some more.

Rich arrays can allow for high-level constructs, and ultimately have all of their procedures deeply coordinated, if indirectly: a connected graph of interrelations ("everything is connected"). Even so, it may consist of many distant procedures (and in particular have a large diameter); cliques that are only thinly connected to other parts of the graph. It may be a product of an extension or amalgamation that contributes only a little (if enough to be worth something sometimes), with few inter-regularities among the original arrays. This makes more room for reducts within the rich and high-level object-type whose value is not subsumed by it. Cats and cars are two *fairly* independent object-types; in most cases, the interactions between these types are not important. Each thus has (in us) its own rather rich, coordinated array of interactions with and within it, along with perceptions of these. However, they may also be part of a far richer object-type that extends them both, and that is a higher object-type too that predicts, for example, the sad consequence of them running into each other.

Maintaining a multitude of partly coinciding object-types creates the added difficulty of choosing the right one. But on the other hand, it allows for the top-down selection of the right "level of description", the object-type that may be relevant to particular goals, while abstracting from the rest. This may thus actually be a better option, computationally, than figuring out predictions and possibilities within an excessively complicated (rich and high-level) object-type. Furthermore (and quite critically), reflecting its place in our mind, the right level of description allows us to *communicate* our intentions and goals to others, directing their focus. At the heart of all these matters is the question of how valuable capturing the extra interactions is in general and in particular circumstances.

I have mapped the basic landscape in terms of the value of some object-types over others, and of some sets of object-types over others. But now, how does the optimization, the search for a better selection of procedure-arrays progress? How does the system actually *learn*? Can it just consider sets of procedures uniformly, detecting and exploring the regularities until it figures out the procedure arrays and which ones it should keep?

Bringing about and managing an object-type is an *orchestration*. The richer and higher-level the object-type is, the harder the orchestration. Figuring out the orchestration of a great multitude of procedures, of rich and high-level object-type, *all at once*, would be highly unlikely (with two caveats to follow). Accepting this entails a fundamentally *bottom-up* view of objectification:

At the beginning, there are only isolated procedures. Hand in hand with the learning of regularities between them, the system strives to bring them together, through extensions, into interrelated arrays, which bring about high-level object-types. New arrays also raise new regularities related to them, which likely could not have been detected of arbitrary composition of isolated procedures (e.g., with regard to the turning of an object, as a new affordance, rather than simply to contractions of unstructured multiplicities of muscles). The system perpetually continues to learn of (and adapt to changes in) regularities between the procedures themselves, as well as of new regularities between them and the new whole arrays that accumulate, and within the new arrays themselves. The system perpetually strives to bring them together, through extensions and amalgamations, into larger interrelated arrays, that bring about higher and higher level object-types (tracking more and more of the structure of reality). This is the fundamental directedness of objectification, the basic acquisition of object-types (which requires the integration of interactions). Thus (for a slogan): *Development recapitulates (the computational-representational grounding of) ontology*.

Against this general directionality, a relatively rich and high-level starting point may in some contexts be possible – and sometimes necessary. Some massive aggregates of procedures (such as the perception from single photoreceptors) may act very similarly, with sweeping *uniform* regularities to their interactions, and so it may be possible (with the aid of innate guidance) to efficiently form for them entire object-types all at once. This may actually be necessary. Too few, separate procedures may not be enough for converging on coherence, which may only arise from a multitude of procedures. For example, just two unrelated photoreceptors behave rather randomly relative to each other (i.e., little mutual information), and so could not be coordinated.

Also against this general bottom-up directionality, there is room for picking out valuable more abstract types too. But let us attend to this while trying to get a taste of the larger picture (pardon the incoherent amalgamation of metaphors); of accumulating real-life object-types.

At some point in the beginning (of a child's life, in the child's mind), it is possible that there are no arrays, no support for ontology of any kind. (If some very-fundamental ones do come innate, it does not matter for this sketch of the general dynamics). Only the objectification system is in place (though it may be still developing, too). The system begins to form a cognitive ontology: object-types such as physical objects (as a general type), food, people, etc. It detects regularities and accordingly constructs richer and higher-level arrays that determine object-types. These regularities can be natural (e.g., just doing stuff moving around the house) or supervised (e.g., being guided by the parents, spanked every time she comes near an electrical outlet); it is all the same to the system (for now, until it can better control its environment and through that, control what it is exposed to, e.g., an interest in animals, or mathematics). The same basic principles drive the system until the end of its life (the focus gradually shifting from amalgamating new arrays to adjusting existing ones, naturally). But certain phenomena are best viewed in the face of it becoming enriched. To tie some former considerations of object-type value to the progression of learning, then:

Amalgamations may be central, and extensions certainly are. But for many amalgamations, and extensions in general, the original arrays should not just be thrown away. Learning about bats should not trample over the object-types of mammals nor winged animals; these are still independently useful (even if there are no other winged mammals). They instead should just be extended the right way (recall our discussion of fruit_g vs. fruit_u, p. 148). Similarly, learning to differentiate between beagles and boxers (and even among all dog types and breeds) does not render useless the lower-level object-type of a *dog*.

However: for object-types that are (approximately) perfectly correlated and interconnected (not having *perceivably* different domains), the original arrays *should* be disposed of. When a mother's voice is always bound to her face, the infant's system should be focused exclusively on the higher-level object that combines both aspects. Perceiving her just as a spatially located face blocks the ability to communicate with her verbally (as mandated by her as a vocal communication object-type). It misses out on the interactions that make the amalgamation worthwhile, and it gains very little in return. There is no point in perceiving her as that object-type, and so its existence in the system is clutter at best.

Ultimately, the procedure arrays that the system keeps as such are the highest level *needed* ones. These as they are used, I postulate, is what the subject is aware of, consciously (in line with the conscious perception of high-level objects in visual perception; Section 5.3). Lower-level arrays that were constituents of kept extensions and amalgamations become automated, and the regularities between them become *integrated*. To no longer bear a cost on conscious attention is precisely, according to the postulation, not to be designated and used by the objectification system as individuated arrays (and accordingly, not have names for them).

In the general learning trajectory, the system is sometimes forced to break down an amalgamation, revive its original arrays, or generally take simple reducts of it, which *turn out* to be independently necessary. Those reducts then *become* (perhaps again) conscious. The aforementioned infant, for example, would later on learn that his mother's face may also appear on a picture, and thus be forced to dissociate it from her voice. This abstraction is a different type of learning than the fundamental (bottom-up) acquisition of novel object-types. The procedures are already there, *coordinated*, and determining an independently useful subset is different (and probably feels different) than learning strictly novel object-types. Such shall be my claim for ordinals and cardinals (Chapter 7).

The extreme point of abstraction is the set of procedures p_O that apply to *all* represented objects; the type-less object-type T_O of simply *object*. This minimal set is devoid of physical procedures, for example (especially if assuming that there are mental objects such as memories). But it need not be completely empty; T_O need not be trivial. Some general *meta-cognitive* procedures are involved in cognitive object manipulation and objectification at large. And there are regularities to them too, which the objectification system may be able to pick up just the same. For example, noticing that (due to some cognitive constraints) there is an informational bound on all object-types (by some perceivable measure) could lead that regularity to become part of what an object *is* (cognitively, metaphysics aside). T_O could then reflect not only the current cognitive state, but to some extent, future extensions of it; something about the notion of an object in the most general cognitive sense.

6.6 Cognitive Assessment

I have presented an abstract framework, which the reader was asked to consider first conceptually, as "just a story of how objectification would or could go". The framework attempts to capture the fundamentals of objectification, as they should (at this level) apply to actual systems, human designed or natural. With most details left open, including entire crucial aspects or sub-systems and parent systems of any such objectification system, the scientific value of the framework is presently limited. More detailed implementations, humanly designed or natural ones, may to some extent be reconceived in terms of the framework. For example, we have seen (Section 6.4) how aspects of empirically explored cognitive systems could be recast in terms of my more abstract (open) notion of object-coherence. However, the framework does frame the topic, committing to some lines of reasoning and speculations along the way, dictating a view that may not square with how the system designers or cognitive scientists may see or frame their own explorations. It does suggest predictions, which reach far beyond the relatively limited empirical input it took for a basis upon which to build (some integrated explicitly within the text, most implicit). How, then, does the framework measure with respect to the relevant disciplines' up-to-date knowledge? How established is the story?

Since the intended application of this framework concerns living beings (first and foremost mathematicians), let us put aside further considerations with regard to robotics, and focus on the cognitive reasonability of the framework. The purpose of this section is to help the cognitively (justifiably) critical reader to see how this abstract framework could be further related to existing, concrete, explorations and findings and understand how it could be evaluated with respect to actual cognitive science reality rather than merely as a coherent story.

6.6.1 Categorization

Let us here attend to a central concrete commitment of the framework, regarding the fundamental directionality of the acquisition of object-types (towards higher-level ones). Unlike various other issues that have been touched upon for the sake of the completeness of the picture, this one is crucial for later on understanding mathematical objectification and using it.

Recall (p. 140) that *categorization* was the familiar topic in cognitive science that in a sense is a (restricted) form of objectification (although our different focus, on much more than just learning to classify objects, generally matters). I have offered an a priori reasoning in support of the fundamental directionality, in terms of how hard it is to orchestrate an array, with respect to how rich and what level height it is. Such reasoning may be enough just to form an expectation; it does not provide any certainty. (Even within the confines of the framework, it could happen to be the case that some basic categories are simply too fundamental with related procedures too interrelated, offering no reasonable sub-array to converge on as an independent object-type along the way). Does categorization in fact develop in this directed way, as suggested? Empirical findings could have produced results to the contrary. In fact, given the topic's classical framing around the *basic level* (which is not necessarily my notion of *highest* level), the prediction was different, and even seemed to find empirical support:

> Rosch et al. (1976) also suggested that basic-level categories might be the first to be acquired during development. Evidence supporting this suggestion came from a task in which 3-year-olds were asked to identify which two of three objects were alike. The key finding was that the children succeeded in the basic-level version of the task involving, for example, two airplanes and a dog, but performed poorly in the superordinate-level version of the task involving, for example, an airplane, a car, and a dog. On this basis, Rosch et al. argued that object categories were initially represented at the basic level. Development consisted of grouping together basic-level representations to form the superordinate level, and differentiating basic-level representations to form the subordinate level. (Quinn, 2010)

However, interpretations of such particular findings do not promise to give the general picture. Despite this early inclination, things have turned the other way:

> [T]he evidence has indicated that global-category representations actually emerge before those at the basic and subordinate levels. (ibid.)

And in some detail:

> The results from across the studies indicate that the trend from more global to more specific category representations is not exclusively found with certain age groups (i.e., older infants) participating in particular types of tasks (i.e., those involving object

> exploration). The broad-to-narrow trajectory is also found for young infants participating in looking-time and ERP tasks. (ibid.)

And in a different methodology, examined through computer simulations rather than empirical research:

> A majority of the simulations produced a common result, namely, that global categories preceded basic-level categories in order of appearance. (ibid.)

A distinction between types of representations which is central to the field, is between *perceptual* and *conceptual* representations. The perceptual is central at early stages, when the infant can hardly act and certainly not learn through language, for example, but mostly just observe. More abstract *concepts*, are attributed to a child in later stages in life. These types are customarily taken to be cognitively disassociated. My framework downplayed differences in representations as non-categorical, through levels and through meta-procedures (that allow for more sophisticated possibilities later on in life). The empirical explorations just discussed turn out to point in a similar direction:

> [P]erceptually based category representations of infants should not be dissociated from the knowledge-rich concepts of adults. Rather, the latter grow out of the former based on a process of enrichment in which category representations that are initially based on perceptual attributes come to incorporate non-obvious attributes acquired through informal and formal tuition, and language. From this perspective, infants begin to form categories based on core competencies consisting of functioning perceptual input systems (including that for language) and a general learning mechanism (that can represent within-category similarity and between-category dissimilarity) that is facilitated in its operation by biases to attend to some inputs more than others". (ibid.)

This suggests a unified framework that conforms with mine, allowing to explore more advanced conceptual learning on par with earlier, first explorations by infants. Another result in similar spirit, expanding the scope of general objectification -

> indicates that cognitive processes like categorization, when put into operation in the developing infant, begin to yield functional knowledge within a short period of time. (ibid.)

To echo yet another central theme of my framework:

> Infant performance is consistent with Malt's (1995) suggestion that many object categories "seem to be strongly influenced by regularities in the input that are recognized by the categorizer" (ibid.)

The large developmental picture is very much like the one I have depicted (Section 6.5):

> If one rejects the notion of a dissociation between perceptual and conceptual, and embraces the notion that category development can be explained via perceptual processes operating in conjunction with a general learning mechanism, then the question arises as to how the category representations of infants transition to the concepts possessed by older children and adults. The studies of infant learning of animal categories suggest that young infants divide the world of objects into perceptual clusters that later come to have conceptual significance for adults (i.e., mammals, cats, tabby cats). As such, the conceptual representations found later in development may be viewed as informational enrichments of the category representations that infants form on the basis of perceptual experience (Quinn & Eimas, 1997). For example, infants who are presented with exemplars of cats and dogs are not experiencing these exemplars as an undifferentiated mass, but rather as separate groups that fall into distinct representations. These representations might then serve as placeholders for the more abstract information that is acquired beyond infancy, through language and learning that occurs in both informal (e.g., home) and formal settings (e.g., school). Thus, over time, the perceptual placeholder representation for cats will come to include information that cats hunt mice, have cat DNA, give birth to kittens, and are labeled as "cats,"… The acquisition of this additional information serves to enrich the original perceptually based category representations to the point that they attain the richness of the more mature conceptual representations of children and adults (Quinn & Eimas, 2000). By this view, what changes as concepts mature is the content of the representations, rather than the processes underlying their development (Madole & Oakes, 1999; Rakison & Poulin - Dubois, 2001). (ibid.)

Though this picture and mine are alike, an important deficiency of *categorization* arises from this latter description. The focus on only the classification (at various levels) of real-world objects, misses out on the finer but critical cognitive distinction, between different object-types that nevertheless share the same real-world ontology[77] (recall the discussion of *fruit*$_g$ vs. *fruit*$_u$, page 148). Any progress between such types is simply outside the scope of this kind of conceptualization, leaving it up to vaguer terms like "informational enrichments". Any connection between such "enrichments" and enrichments that through distinctions also break apart the category into subcategories, is lost upon the system. An example of this – is the proposed fundamental principle towards the higher-level (Section 6.5), as a unifying, general principle.

6.6.2 The Two Streams Hypothesis

Following Gibson's focus on perception-for-action (that my framework has integrated into the object-types), arose (about two decades ago) a neuroscience hypothesis, which suggests that this is supported by a second system, separate from what is usually meant by "perception". Since then this hypothesis has been finding much support in the research and gaining much ground. It is important to note the cognitive complexities that may threaten my framework. To start with a particularly-contrary presentation of the issue:

> One clear fundamental notion of the affordance concept is that object recognition is not a necessary step for interacting with objects. That is, a specific combination of object properties with respect to the agent and its action capabilities are enough to detect the affordances of a given object (and act on it). Although it is not the classical engineering approach of identify and then act, this strategy appears to be the one employed by our brains. It is known that the cerebral cortex processes visual information in at least two channels, the so-called dorsal and ventral pathways [or "streams"]. The ventral pathway appears to be responsible for object identification, whereas the dorsal pathway is mainly involved in perception for action [18–21].[78]

---

[77] For another exemplar, take: "this alternative account suggests that the early development of categorization may proceed from more general to more specific representations based on a single system of representation that becomes progressively differentiated during the course of experience" (Quinn, 2010) (my emphasis).

[78] A bit more elaborately:
"They argued there are two visual systems, each fulfilling a different function. First, there is a vision-for-perception system based on the ventral pathway… which is the one we immediately think of when considering visual perception. It is the system we use to decide that the animal in front of us is a cat or a

Factor	Ventral system	Dorsal system
1. Function	Recognition/identification	Visually guided behaviour
2. Sensitivity	High spatial frequencies: details	High temporal frequencies: motion
3. Memory	Memory-based (stored representations)	Only very short-term storage
4. Speed	Relatively slow	Relatively fast
5. Consciousness	Typically high	Typically low
6. Frame of reference	Allocentric or object-centred	Egocentric or body-centred
7. Visual input	Mainly foveal or parafoveal	Across retina
8. Monocular vision	Generally reasonably small effects	Often large effects (e.g., motion parallax)

Fig. 11 Eight main differences between the ventral and dorsal systems *(Eysenck, et al., 2010 p. 48)*

> These data suggest that an agent does not necessarily need to possess object recognition capabilities to learn about its environment, and use this knowledge for making plans. (Ugur, et al., 2011)

Our main focus has been on object-*types*, and for determining just a type, "a specific combination of object properties" might suffice, particular object recognition aside. But this would certainly go against the spirit of my approach, which strived to relate perception and action to the object-representation, as such, identified, full-blown. If object *representation* (never mind precise identification) turns out not to be central to perception and action, at least at the top-most, conscious levels, then there is much less to build on for understanding the representations of mathematical objects through the cognition behind the representations of standard objects (that may be innately designed for). Cognitive presentations may refrain from such strong conceptual conclusions as above, but this separation, for which finding a place within my own framework would not be easy, are certainly substantial:

But the empirical explorers are still hard at work, conceptualizations change, and in particular explorations of the *connections*[79] between the two "streams" come to defy their earlier, simplistic classifications:

---

buffalo or to admire a magnificent landscape. In other words, it is used to identify objects. Second, there is a vision-for-action system (based on the dorsal pathway…, which is used for visually guided action. It is the system we use when running to return a ball at tennis or some other sport. It is also the system we use when grasping an object." (Eysenck, et al., 2010 p. 47)

[79] For example:

(1) The ventral or “what” pathway that culminates in the inferotemporal cortex is mainly concerned with form and colour processing, whereas the dorsal (“where” or “how”) pathway culminating in the parietal cortex is more concerned with movement processing.

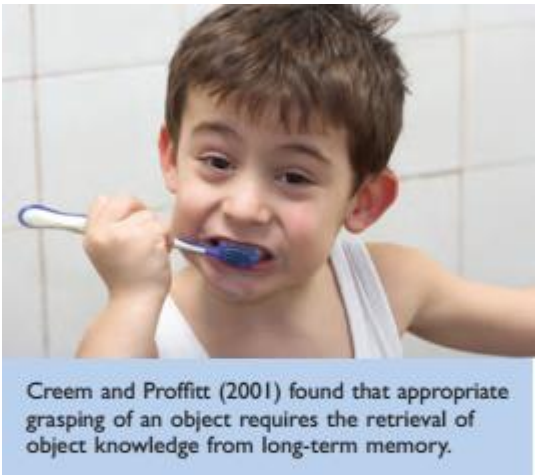

Fig. 12 Differences between effective and appropriate grasping *(Eysenck, et al., 2010 p. 54)*.

(2) There is by no means an absolutely rigid distinction between the types of information processed by the two streams. For example, Gur and Snodderly (2007) discovered a pathway by which motion relevant information reaches the ventral stream directly without involving the dorsal stream.

(3) The two pathways are not totally segregated. There are many interconnections between the ventral and dorsal pathways or streams. For example, both streams project to the primary motor cortex (Rossetti & Pisella, 2002). (Eysenck, et al., 2010 p. 38)

Just as an example: Grasping a physical object (and perceiving the ways it could be grasped), is partly a “core” ability, supported by the dorsal stream (possibly avoiding a full-blown object-representation, seeing just the grasps it affords). However, since this system does not retrieve representations stored in long-term memory, there must be more to grasping than that, as we can gain and master knowledge of grasping particular object-types (e.g. chopsticks). This caveat calls for a distinction between (what has been termed) “effective” and “appropriate” grasping (Fig. 12). Such supposed interplays between the two systems can be quite intricate.

Finding a path within this complex topic is far beyond our scope here. But this example hopefully gives a taste of how my account relates to debates and tensions within cognitive science.

“because dorsal and ventral visual regions are heavily interconnected, it is difficult to tell in healthy subjects whether information is processed along the dorsal stream only, or whether it is fed to parietal cortex via ventral visual regions”. (Hebart, et al., 2012)

6.7 Mathematics

Explicit mathematics in general, and its types of objects in particular, by and large need to be learned. Some very basic and highly applicable bits, such as quantitative comparisons, may be fundamental enough in real life to possibly earn innateness. But even this is not at all certain: being fundamental and abundant means that they could instead also be learned, where all that needs to be innate is the ability to pick them up from the environment. This might be true for quantities if not numbers, for example (to contrast with Dehaene's innate intuitions grounded view; Section 3.2). But numbers, furthermore, seem to require at least some non-innate learning, as they are dependent not only on sufficient age (which in itself does not rule out innateness), but also guidance and practice. To the extent that mathematicians do come to learn any *new* types of objects, the capacity to do so is part of the full cognitive system. And ruling out a fitting adaptation for mathematics at the evolutionary scale, this capacity is most likely a *general* one, ranging over various non-innate, non-necessary realms besides mathematics. I.e. – the objectification system.

Mathematics' objects are thus suggested to be on par with regular, non-mathematical objects, through how we come to *attribute* objecthood in the other realms. As discussed previously (Section 4.2), this seems to me to be the right default route, until differences that arise (as they surely will) force us to split our understanding of the cognition underlying mathematics from the rest.

Approaching the cognitive constitution of mathematical objects, along the lines of the general framework provided, is my offering of (first steps towards) a *representational foundation* (Section 1.5). Metaphysically, it could serve as a first cognitive-metaphysical step on way to a cognitive foundation (Section 1.6 & Chapter 2), then grounding mathematical objects in their object-representations, perhaps. But in general, I strive to avoid making additional metaphysical commitments. The framework is used in what follows only for exploring mathematical ontology through its cognitive representation.

Let us now conclude the account of the objectification framework with two issues that are directly mathematical.

6.7.1 Regularities

Regularities (Section 6.4) are key to the structure of ontology, and to its tracking through extensions and amalgamations. In general, among all of the different kinds of grasped regularities (whose metaphysical type the system does not discern), some that are empirically evident in the world are actually *regularities of the cognitive system*. Some of these are temporary (as an infant's vision stops being blurry, for example), while others are more stable and long lasting. Some are fixed for any human cognition – and some even for any cognition at all. It is the latter that I posit are the mathematical regularities – *metaphysically*. But to the extent that mathematics does share this *general* objectification framework, there is no a priori reason to attribute a special *cognitive* status to approximately perfect regularities that happen to be *mathematical* (and particularly perfect). For matters of objectification, the system can already perceive them as regularities. Perceiving them *as* mathematical is a different story altogether, one that is not grounded in statistically based learning, and for which the framework does not presume to account. (Doing so may be beyond the reach of current cognitive science and artificial intelligence (Sloman, 2016)).

Having all approximately perfect types of regularities meshed together has cautionary implications for our perception of mathematical objects. For some perceived mathematical object-types, the object-coherence of the supporting arrays may depend on non-mathematical or even non-necessary regularities; Chapter 7 demonstrates precisely this.

6.7.2 Structures

The objectification system acquires and manages cognitive object-types. Having an object-type can bring with it a *schematic* conception of what it is like for the objectification system to 1) recognize and perceive a given object as an object of that type and 2) activate one of the available procedures on it. In this sense, one can know, to a pragmatically acceptable degree, what apples are, for instance; “I will know it when I see it” and “If it is an apple, then I can bite it, or see if it is rotten”. This is a relatively simple meta-cognitive reflection upon the basic operation of the objectification system, *with respect to* a given object-type. It is still a form of abstract thinking with respect to the object-*type* rather than just an actual, particular perceived object (and so perhaps some creatures that do objectify may not be capable of this type of consideration). In itself, though, it is *not* a reflection upon the *interaction* between the whole *array* of procedures that support the object-type; just upon one, or the interaction between a few ("Sure, if it's an apple then I can see if it's rotten, and if it isn't, I'll bite it").

The full object-type, as its whole array of procedures, may furthermore *determine* the category of objects (in the world) to which the object-type applies (though following Rosch et al., the boundaries may ordinarily be graded).[80] But the category, as a whole, is (by default) nowhere to be found represented *in* the objectification system. For an order-type, the system does not manage its domain as such; this aspect is not part of its operation, not part of what it is set for (and not part of ordinary life). One can have a schematic conception without having any clear conception of the set of the apples in the world, the set of all possible apples, or any idea of that sort. One could be surprised, in a sense, by a black apple, without the objectification system's rejecting of it as an apple altogether. The objectification system has simply never attended to the issue before and has no stance on it, no representation of it (and nor does the larger system, the person reflecting upon the objectification system, perhaps). This is despite the fact that it was *theoretically determined* for the object-type of *apple* that it *would* apply to a black apple too.

Grasping a *structure*, the whole domain of objects of an object-type and the relations between them, is a hallmark of mathematical thinking. It is a product of a deeper meta-cognitive reflection, from outside of the objectification system, upon the *full* object-type as it sits within the objectification system. Even for a *mathematical* object-type acquired as such, such structure grasp does not come by default. It requires reflecting upon the array of interacting procedures just as well, not merely knowing when and how to use them and reasoning *locally* with regard to a few. This further holistic, sophisticated reflection may be challenging or even infeasible.

Couched in these abstract terms, why would this holistic reflection be a hallmark of only mathematical thinking? I suggest that this sort of achievement is simply not available for non-mathematical object-types. These actually do depend on the world. The reflection upon them reflects that, and so does not presume to capture ontology

[80] If it actually does, is a metaphysical question that we can avoid for now.

through cognitive representation alone. Furthermore, such object-types are far richer in terms of the procedures they make use of, which involve innate, hidden mechanisms or just far deeper automated support. For structure grasp, the procedures need to be relatively accessible to the reflection. This activity can thus begin when we *abstract* from the world (and the rich mechanisms that support its perception).

All of this certainly requires a further deep account – which, for our particular purposes, is not needed. However, the distinction itself, between merely acquiring an object-type and grasping its structure, should be kept in mind as we turn to mathematics.

7 Ordinals vs. Cardinals in ℕ and Beyond

In Chapter 6, I presented a general framework for cognitive object-constitution. Here, I turn to the object-types of ordinals, cardinals, and numbers. It is up to the objectification system to construct these types, too, by learning the relevant regularities and converging on these object-types' arrays.[81] The focus is on the intricate *relation among the three object-types*, in between the finite and the infinite.

I use the *developmental* framework (accounting for the mastery of ℕ in early childhood) to unveil an *ontological* picture. It is different from the familiar related philosophical accounts (of cardinals vs. ordinals), and disputes the classic mathematical tale of what actual infinity supposedly did to our concept of numbers.

7.1 The Classical Mathematical Narrative

The legitimacy of actual infinity in mathematics has been a central issue of controversy, from the very early days of Aristotle (who distinguished it from potential infinity) and all the way through to Gauss and beyond. What could infinite numbers be taken to be? They seemed to be inherently paradoxical, contradicting fundamental intuitions. Finally, Cantor legitimized them by showing that they can actually be handled coherently, to a mathematical degree, if one can give up some of those intuitions.

The famous aspect of Cantor's move was to base equicardinality, one central aspect of numbers and their usage, on 1-1 onto and total mappings alone (where being a subset no longer entailed being strictly smaller). Before that, in the finite, basing equicardinality on mappings alone was *possible*. But so was basing equicardinality on set containment. It made no difference; the result was the same because the two approaches were *synchronized*: one could safely switch – in mid-air (i.e. mid-reasoning) – between subsethood and mappinghood (the nonexistence of a surjective mapping). It is only in the face of actual infinity, so this classical narrative says, that the notion of quantity breaks, and one has to give up a central part of it.

[81] See (Sloman, 1978), Chapter 8, for example, for a rather thorough list of computational requirements for being able to grasp numbers.

Cantor did not develop only this up to 1-1 correspondence side of numbers as taken into infinity. When arbitrary 1-1 mappings are disallowed, yet another aspect of numbers can be extended to infinity: that of (well-)orderings. Ordinality too reflects a basic aspect of numbers and their usage. And it, too, can be made mathematically coherent and useful in the infinite too.

Beyond the finite, ordinals and cardinals simply behave differently – incompatibly so. In particular, the same object cannot always be considered both as a cardinal and as an ordinal. The notions split, and the ontology itself breaks up. The classical narrative in mathematics, then, says that when taken to infinity, the concept of numbers must be broken between standing for quantities and lengths. But in the finite, numbers are the fundamental entity. They may be given to us by God (as Kronecker would have them); or modernly, given through a logical axiomatization, such as (second-order) Peano Arithmetic. Once determined as a structure, that is as far as their mathematical (Sui generis) essence goes. They may be *used* ordinally, or cardinally, but that is just a matter of the *application* of mathematics (expressed in various ways in the cognitive, linguistic, or developmental arena).

Or so the classical mathematical narrative goes.

7.2 In the Finite

The infinite forced the acknowledgement that ordinals and cardinals are different types of entities, each with their own, incompatible arithmetic. The full extent of this difference is obscured in two ways in formalized set theory:

i) A common trick reframes cardinals as "a type of" ordinals. Since each ordinal, as a set, has some cardinality, ordinals can be used to represent cardinals. Since ordinals are well ordered, it is well-defined to take the smallest one of each cardinality to represent that cardinality. However, being well defined and even canonical does not make it *natural*. There is no (metaphysical or cognitive) reason to privilege the first ordinal. Moreover, the relation of *being a representative for* another entity is an admission that there is something extra going on (some computation or reinterpretation); that an ordinal and its cardinal are not just one and the same. $\omega = \aleph_0$ is not a logical triviality.

ii) That every cardinality can be represented by an ordinal, is the axiom of well ordering, more (in)famously known through its equivalent, as the axiom of choice (AC).

To the extent that it is coherent to even consider a world of sets in which AC does not hold, then the general notion of cardinality is thus dissociated from ordinality. Cardinality cannot even be built upon ordinality.

And of course, the *von Neumann* ordinals themselves are just a form of representation through sets. For my project, we need to put all such coding tricks aside. They may suffice for the proof-theoretic outlook, but not when we seek the concepts themselves, as they appear to us and underneath that. Thus, let us roll back from ordinals and cardinals *as sets*, to a notion more similar to Cantor's earlier conceptions of equivalence classes (of being order-isomorphic and set-isomorphic, respectively). Having an underlying set-theoretic semantics is not of our concern here (and our focus is on the finite anyway).

Ordinals and cardinals, then, really are inherently distinct objects; a *mathematical* difference rather than just a matter of application. Even if this becomes evident through the infinite, should it not apply to the finite ones just as well? Might they too be brought together by extra connections that hide their differences and manage them, into what we perceive simply as natural numbers? Is this a meaningful question, when we are no longer *forced* to differentiate them? Might differences between them still express themselves in the finite somehow? Steiner (2005) (for a modern presentation) shows that this is indeed the case.

> [T]he [finite] cardinal–ordinal *equivalence* has another implication, and a profound one. This is that cardinal addition is equivalent to ordinal addition; extensionally, they are the same operation. But ordinal addition is based on recursion (iteration), and recursion yields computational algorithms, such as the ones we learn in school… these algorithms are for ordinal arithmetic, which has little to do with the more applicable cardinal arithmetic, which measures the sizes of various sets… arithmetic can be immediately applied to cardinal arithmetic to make calculation a cinch. The mathematicians suffer from the lack of algorithms in infinite cardinal arithmetic, and as a result cannot make the simplest calculations, such as $2^{\aleph_0}$.
>
> Cardinal arithmetic, then, is of great application but does not allow calculation; ordinal arithmetic has no great application, but does allow algorithmic calculation. Together we have the following story: take two distinct collection – of bodies – and count them both. We then have the ordinal number, and thus the cardinal number, of the two

> sets corresponding to the collections. The cardinal number of the (disjoint) union of the sets corresponds to the sum of the two cardinal numbers. This can be calculated using ordinal arithmetic (i.e., algorithms). From this we can predict the result of counting the… sum of the two collections.

It should be stressed here that *counting* itself mixes the two: to "point to all the bodies in a collection while reciting numerals in order" is to assign a particular order, i.e. ordinally. But counting is inherently order independent. (One would not correct another if she opted for a different ordering than the asker had in mind). Counting operates ordinally to answer the question of "how many", cardinally. That the final number-word in a count sequence, which is attributed to the last object counted, indicates the cardinal value of the *set*, is known (in the developmental mathematical cognition literature) as the *Cardinal Principle*.

The arithmetic operations are defined, built, through different constructions. Order dependence or independence must be integrated into them. Importantly (as is discussed in the following sections), an order-type is *not* a set *together with* an ordering of its elements, but just the ordering of elements itself; an *underlying symmetrized* conception need not be involved. Two orderings of objects, *a* and *b* (instances of ordinals), then, may each be turned into the matching symmetrized (unordered) true sets (instances of cardinals). But concatenating them, the way ordinals are added, would produce only a partly ordered "semi-set" in which all of *a*'s objects, in whatever order, come before all of *b*'s objects, which are in whatever order as well. Instead, a (disjoint) union must be used for addition, an operation directly from (symmetrized) sets to (symmetrized) sets. Conversely, ordinal addition cannot pass through set union, or else the order would be lost. And similarly for all (of each type's) arithmetic operations. Accordingly, (Steiner) goes on to demonstrate why the two notions should be kept separate:

> Why is it always the case that adding x, y times, gives the same result as adding y, x times?... The "definition" of multiplication as repeated addition has its source in ordinal arithmetic (i.e., recursion)… Small wonder that children cannot understand why multiplication is commutative (the noted philosopher and logician Saul Kripke recalled… being amazed as a child at the commutativity of multiplication). Nor can they understand why one can "add candies to candies" but not "multiply candies by candies". What we need is a **cardinal** definition of multiplication…

> A good way to look at multiplication is, in fact, set-theoretic: xy is the (cardinal) number of x disjoint sets, each of which has the cardinal number y… we see from this definition that each multiplication is equivalent to (though not synonymous with) a repeated addition—we can "read off" the results of a repeated addition from a product of two numbers.

Here, x and y still play a different role (which is why one cannot "multiply candies by candies"). But this cardinal analogue of repeated addition (Steiner continues) has a 1-1 mapping to the Cartesian product x×y. Furthermore, x×y is in a 1-1 mapping with y×x (pictorially, by e.g. rotating the rectangle by 90°). A set of three candies assigned to each of four children could instead be assigned by fours to each of three children. Thus, x and y play a similar role, and cardinal multiplication is easily seen to be symmetric (even if not defined directly through the Cartesian product). To summarize, seeing that counting to x (as an activity of ordering) y times, would result the same as counting to y x times, passes through cardinal intermediary stages:

counting to x, y times ⇒
y disjoint sets of cardinality x ⇒
x×y ⇒
y×x ⇒
x disjoint sets of cardinality y ⇒
counting to y, x times

This transition, of course, cannot be done beyond the finite, where ordinal multiplication indeed fails to always be commutative. Understanding the practice of counting-qua-ordering alone should *not* lead to seeing the commutativity of its arithmetic. (Ordinal multiplicativity can alternatively be proven by induction, but 1) it is doubtful that children can come to see commutativity through induction, even implicitly, and 2) induction – strictly finite induction – presupposes a grasp of the finite).

Numbers-qua-ordinals versus numbers-qua-cardinals, and how they are differently used, differently *applied*, brings up Frege's constraint, which is, "roughly, the requirement that the core empirical applications for a class of numbers be "built directly into" their formal characterization" (Snyder, et al., 2018). Ordinals and cardinals are different – and indeed so are their "most general principles governing the empirical applications". (And what about just *numbers*, then?)

That this sort of analysis (of Steiner's) can exist, revealing finer workings within a seemingly extremely familiar and basic realm, shows that something is missing from explicit foundations that take numbers as a given, or as given through the laws they uphold; and that perhaps we were *consciously* missing this thing (mathematically fundamental) that our brains have nonetheless learned to manage successfully and systematically. The implication is that for our exploratory purposes for foundations, we must dive into the cognition that hid these workings. *The objectification framework I have developed serves as the framework in which such analyses and explorations can occur.*

7.3 Philosophical Narratives

Natural numbers (as even their name suggests) might ordinarily be introspectively perceived as *sui generis* objects, that have various properties but that cannot be reduced to any more fundamental entities. Foundations of mathematics started out by transcending this appearance (Section 1.1), striving to ground them in logic. Centrally, ordinality and cardinality were two different aspects of numbers, and their place within the nature of numbers has been a topic of some debate. Most of it has concerned the *priority between* these aspects (Samuels, et al., 2017). I very briefly mention a few central figures here in order to locate the account of this chapter within this debate. (Taking numbers as real objects, nominalism, fictionalism, etc. are not mentioned).

Frege's account of cardinal and natural numbers[82] is probably the most well known or paradigmatic. It takes cardinality to be fundamental, *cardinal* numbers constituted as independent classes of equinumerous concepts (or sets). An ordering is then imposed on them. This is done using the (immediate) *predecessor* as a *cardinal* relation; namely, having one less object, generalized purely *ordinally* into the ancestral relation, of being the predecessor of the predecessor of … any finite number of times (but without any reference to *number*). The natural numbers, as a unified concept (rather than separate

[82] For modern presentations, in terms of terminology and with regard to avoiding the inconsistency of his Basic Law V, see (Zalta, 2017; Demopoulos, et al., 2005).

objects of cardinality), can then be defined as being 0_{Card} itself or having 0_{Card} for an ancestor.

Cardinal numbers are completely general here. But the particular order imposed on them comes from the finiteness of the logic and tracks it[83]. More importantly, there are no ordinal numbers, explicitly, as an independent type of object. Cardinality is given an essential priority. In the face of Cantor's development (at around the same time) of both infinite ordinals and infinite cardinals (clarifying the fundamental distinction), this approach now seems less reasonable.[84]

A critical feature of Frege's account (following Chapter 1) was its aspiration to be not only an *epistemic foundation* but a *metaphysical foundation*, one that accounts for what natural numbers really are (in his philosophy, through their grounding in logic). It was not just a *mapping* of natural numbers as a concept – whatever and wherever they may be – to a logical, rigorous arena through which we can come to know them and in which we may find security. This aspiration brought about the problem of the need to tell a number apart from, for example, Julius Caesar (Zalta, 2017). This aspiration, in particular, meant getting the relation between ordinality and cardinality correct and explicit. No tricks like set theory's codification (previous section) would suffice.

Russell followed a similar, cardinality-based path (and the above definition of cardinality is named after them together).[85] For example, after explaining how, in counting, we first order and only then deduce the cardinality from it, he notes the following:

> But this result, besides being only applicable to finite collections, depends upon and assumes the fact that two classes which are similar [equinumerous] have the same number of terms; for what we do when we count (say) 10 objects is to show that the set of these objects is similar to the set of numbers 1 to 10. The notion of similarity is logically presupposed in the operation of counting, and is logically simpler though less familiar. In counting, it is necessary to take the objects counted in a certain order,

---

[83] Though perhaps it could be generalized.

[84] For the tension between Frege's concept of number and Cantor and Dedekind's, see (Tait, 1996).

[85] Though upon a very different basis, using types to avoid Frege's Basic Law V. (Demopoulos, et al., 2005 p. 135).

as first, second, third, etc., but order is not of the essence of number: it is an irrelevant addition, an unnecessary complication from the logical point of view. The notion of similarity does not demand order. (Russell, 1919 p. 17)

Russell is correct to point out the dependence upon *similarity* (which is worth keeping in mind in this dissertation, as it applies to the similarity between objects and *their representations*, too). However, similarity, or *isomorphism*, is a more general (namely category theoretic) notion than the set-isomorphism (i.e. bijection) that Russell has in mind here. The ordering of objects according to numerals is an *order*-isomorphism, its own kind of isomorphism, which differs from that "similarity" of having *some set* isomorphism (even if it happens to be a *type* of the latter).[86]

His point that the necessitated temporal order of things is not the only, and perhaps not the most important sense of priority, is an interesting and important one. But the independence of ordinality from cardinality is more substantial than that. As mentioned in the previous section, it is standard in formalized set theory to even consider it to be the other way around, in a sense. Stemming from a very different philosophical point of view, phenomenological rather than logical, Brouwer has prioritized ordinals even in finite numbers, as constructions:

Discrete mathematics begins for Brouwer with the operation of forming ordered pairs of distinguished elements, and continues by considering repeated iterations of that activity. This produces finite mathematics. For such iterations, says Brouwer, generate each natural number (Posy, 2005 p. 323)

Many years later, Dummett (1991) took the independence of the result of counting from its order to show that (p. 53) Frege has been correct to take cardinality as intrinsic to the characterization of natural numbers. Later on (p. 293) he took the same fact (besides the insights from Cantor's work) to demonstrate that the natural numbers should have been characterized as finite ordinals instead (Snyder, et al., 2018 pp. 87-88). [87]

Even quite recently, for Linnebo, "The views found in the literature *naturally* fall into two types as follows: those that take the natural numbers to be individuated as cardinal

[86] Something similar seems to be suggested by (Samuels, et al., 2017).
[87] A seeming contradiction pointed out by Samuels et al. (2017)

numbers, and those that take them to be individuated as ordinal numbers" (2009 p. 223) (my emphasis). Under this false dichotomist framing, he then takes the various objections he raises to the cardinal-individuation view as an argument *for* his ordinal-individuation view (and his own account along those lines).

The point in all of this is that these are delicate issues, wherein priority between the types is but one aspect within the general question of what the natural numbers are; and that there may be different senses of priority at play (which Snyder et al. (2018) try to map). The Logicist point of view prioritized certain, more abstract senses of priority, whereas Brouwer's phenomenological one stressed a creative, iterated sense. From our cognitive foundations point of view, priorities within the representational (computational) structure and along the developmental track (where these two are not necessarily the same) cannot be ignored. It would take a strong metaphysical argument to refute their relevance and reverse the picture they suggest. At least as long as (as in Chapter 2) there are no better alternatives to build on (with logic no longer qualifying either).

The account I offer here follows the bottom-up formal-developmental lead of the objectification framework (of Chapter 5). Although its focus is on the interaction between the finite and the infinite, along the way, it supplies a blueprint for the natural numbers that diverges from the existing approaches.

7.4 Ordinals

As previously discussed (Sections 6.4 & 6.7.1), not all procedures and their regularities are directly about the physical world. Putting entities in order and perceiving them as ordered is one area of activity governed by its own regularities that *could* be considered *cognitive* (following the discussion there). It can be spatial (putting different toy-soldiers in a row); or temporal (put spoon in food, put spoon in mouth, swallow, which can be concatenated to itself, or, when no longer hungry, to throw food on floor and then parent comes to clean); and it can be non-physical (recalling putting the toy-soldiers in a row); or even a hybrid (light goes out, parent leaves, feel scared, try to think of ice-cream, imagine a monster instead, imagine a scarier one, cry, parent comes back). It can be about anything that the mind can take for (individuated) objects, stripped of any extra content.

Metaphysically, there is no need to stop there. One could still believe furthermore in some cognition-independent platonic reality in which the related ordinal objects exist, whereby cognition only (ideally) reflects or computationally embodies their mathematical properties. But our business here is in accounting for mathematical objects and ontology as cognitive object-types. All that is important for this is that the mathematics involved *also* reigns over the patterns *as* cognitive ones.

Back to the cognitive, then. To be sure, orderings and ordinal regularities are about much more than putting a bunch of stuff in order. For example: A song which is to be sung has three verses, say. They can each be decomposed into four lines. Reframing the list of verses to sing as four mini-activities for each verse should produce a list of 12 total such mini-activities. This cognitive reframing, grasped abstractly as a procedural manipulation of ordinal multiplication, should come to be coordinated with the agent's perception of the order.

Such regularities can reign over all object realms. In particular, a meta-regularity about the system itself presents itself: as new object-types are learned, they are seen to also fall under these ordering regularities. This type of generality is what abstracts away from the particulars of each and every specific realm, including potential future ones. All that is left is ordering's ability to pertain over (a multitude of) the system's most general notion of an object itself (p. 170). Some blobs of color could be perceived as an apple. Two people could sometimes be perceived as a romantic couple. Likewise, perceived configurations of *any* objects that are ordered configurations (or given through an ordering process), could in principle be perceived abstractly as an order-type – and acted upon as such, the content of the object place-holders aside. (Much more should be said on the place of abstract, mathematical objects in the cognitive system. But for the more specific purposes of this chapter, this is sufficient).

The system, in performing its function, can come to detect the regularities between the meta-cognitive procedures involved in ordering-handling. It coherentizes them, adjusts them into coherence, ideally converging on a procedure array P_{Ord} that brings about the object-type T_{Ord} of being an ordinal. In early childhood (the life stage we consider here), this convergence need not stop precisely there. For one, consider the regularity of being *well*-ordered: As any ordering in the finite is a well-ordering, this regularity too, if detected, may be integrated into the array. However, noticing this sophisticated property (intimately tied with *induction*) is far from trivial, and it comes much later than some simpler, more general array of *ordering* probably does, while building upon it (the way the system enriches object-types). Another fundamental feature that early mastery might lack is that of being an order-*type*. The early related notion seems somewhat different, as it denotes the place of single elements *within* the structure (and privileging the last item). The holistic order-*type* may be but a theoretical construct that is not part of the actual array at this stage. There certainly is interest in exploring such differences (from full-blown ordinality), even in the finite, but this is beside the point of this case study, and so I leave it aside. Much more importantly for this story, the objectification could in principle land directly on something larger, namely ℕ, never coming to individuate anything like the object-type T_{Ord} or some reduct. This possibility is not theoretically reasonable nor empirically evident, but we dive into these developmental complexities in Section 7.7. For now, our focus is on the cognitive object-type itself.

Following Section 6.7.2 on structures: A possible theoretical construct that surely is *not* part of the array is *the ordinals*, as a structure. Not the proper class, not some initial infinite segment, nor the finite one, by default. The holistic conception of a delimited domain of ordinals does not come with the mere acquisition of the object-type and the ability to recognize an object as of that type, schematically. Rather, it requires reflecting upon the whole of P_{Ord}, not merely knowing when and how to use it, nor even reflecting upon local interactions between some of its procedures.

With the domain itself not perceived in any such sense, not delimited, a central issue (for us) is left void: could there be infinite ordinals? As all of the concrete ordinals that have been presented to the child have in fact been finite, she might generalize conservatively (as a general rule-of-thumb for cognitive generalization), to similarly finite new ones. She may not be surprised by these and may know, on principle (at some stage), that they can be arrived at through iterations of the successor. Ruling out the possibility of any other ordinals requires more than that, though: it requires the holistic grasp of all that is P_{Ord} and what it allows for. The central distinction here is between all that P_{Ord} or a reduct of it can *generate* and all that it can *operate on coherently*. Generation would be restricted early on to the finite, before the child has acquired, for example, the meta-procedure of creating a position that follows all that the successor procedure ever generates from 1 (namely, ω). Even this step requires holistic reflection that is not part of acquiring P_{Ord} itself (though it is rather simple in this case, as any other acquired operation, e.g. multiplication, can independently be seen to generate nothing new). It is the latter determined domain, of all that P_{Ord} can operate on coherently, that is *the domain* of its object-type. It supports possible procedures that have not been acquired (and it is much harder to figure out than the former generatable domain). For example, consider how the mentioned meta-procedure that generates ω may be introduced: The child is presented with the Zeno paradox in which, before a runner could reach the end of the path, he must reach half of the distance to it. The child may certainly be baffled by the question of how the runner could ever make it to the end of the path. But what *is* coherently perceived is the suggested ordering in the background, of the infinity of mid-way points *followed* by the endpoint. ω can quite instantaneously be recognized as a valid (albeit novel) object of T_{Ord}. All this knowledge comes later. But at the stage of the successful acquisition of just T_{Ord} itself, whether there are or could be *infinite* ordinals is simply not an issue yet reflected *in* the system. Not before introducing this further concept or distinction or, later on, infinite exemplars and constructions.

To add one final point, concerning T_{Ord}'s distinctness from cardinality (the topic of the next section), recall (section 7.2) that we approach ordinals as order-types rather than through set-representatives. In particular, an order-type in itself isn't a set. But more interestingly, an *ordering of objects* (a concrete manifestation of the ordinal, its application to reality) actually need not cognitively be an *ordering of the set* of those objects. The set of the ordered elements is a *theoretical construct*, that can be deduced (by grouping the elements into a set) but probably isn't part of the array itself. P_{Ord} alone can to some extent function independently of any symmetrized notion of set. That is, it may be possible to have T_{Ord} grasped by the system, without having any ability to perceive permutations as in any sense the same. Ordinals can still be concatenated etc. And in cases where the cognitive system does entertain some objects in an orderly fashion, it may grasp the setting abstractly, abstract its order-type from the rest of its contents. For example: It may have the ability to recognize an ant. And seeing a line of walking ants, it may perceive one and then, following the line, perceive the others as ordered appropriately. Perceiving the ant-line as a whole *set of ants*, independent of the order in which they were presented, need not be involved. Where sets do come in is in the particular activity of *ordering* – i.e. of a pre-existing, a priori symmetric set. But this takes place in a wider context than just ordinals, which we will get to soon enough.

7.5 Cardinals

Perceiving cardinality (and in particular perceiving a permutated ordering as *the same*), perceiving just what is invariant under permutations and having procedures operate on *that*, means a *different object-type*. The procedures are not identical and are not used identically. The relevant regularities need to be detected by the system and be made part of the object-type if it is to be object-coherent. For example, for a permutation followed by another permutation, there is always a single permutation that results in the same permuting (i.e. compositionality). Without this, without the transitivity of object-"identity" (as decreed by this object-type), objecthood would collapse. (And similarly for the existence of an inverse permutation). Assimilating these regularities, we can then come to see, for example, that multiplication is commutative, as it is a type of permutation (where the axes are switched). Or more concretely, that turning a table would not change how many pens there are on it (counting always from left to right).

Let us compare the case to an example of physical objects. Such objects (standard, in three dimensions) can be rotated (on two dimensions). Rotating them ordinarily has some effect that may be useful (e.g., turning the chair to face the TV). And coordinately, their being or having been rotated can be perceived. There is, however, a certain type of object that is (in central respects) indifferent to rotations: balls. Invariance under this rigid type of transformations means that balls obey symmetries similar to sets-over-orderings (if less general, including only *some* point isomorphisms). The angle of a ball (towards us or towards the wall, say) can be constructed, it still “exists”, just as it does for other objects. The added symmetry may seem to mean that the object-type is not as rich (having no substantially different orientations to be in). But this is a regularity, which by now we know can also allow for a higher-level, enriched object-type. For example, a ball can be thrown at a wall – coordinated (after some practice) with catching it, in whatever angle it is in the hand and at whatever angle it hits the wall. Ideally, a less symmetric shape would simply allow for more possibilities by turning it in different ways and thus making it bounce in various directions. But cognition is of course constrained – in our capacity to attend to the object's angle, to calculate the direction it would bounce to, and to master the ability to actually make it hit a wall at a particular angle. Symmetry can be of value. Thus, while in the worldly sense, balls may be just a *certain type of* physical objects, not a *different* type of object, representing the ball's angle is meaningless, and hence pointless[88]. The higher-level symmetrized cognitive object-type is of more value.

In cardinality, which is not merely a type of ordering in which the order does not matter (even if they could be coded as such), this symmetrical representation is inherent. The abstraction rids the object-type of every aspect that does not make a difference (to echo Benacerraf's point that numbers are not sets, in any particular coding). An ordering can be reordered (possibly producing a different order-type in the infinite), while a set cannot, as it has no order. The differences between ordinals and cardinals express themselves even in their most basic (and in particular finite) applications, as we have already seen (Section 7.2).

All of this is not to say that ordinals and cardinals do not share a lot, cognitively. They may certainly share many procedures, which can belong to more than a single procedure array. Moreover, one may in some sense actually be "built" upon the other. Perhaps cardinality is understood through grasping the symmetrization of ordinals (an invariance of action and perception). The procedure of counting (cardinality through ordering) and the grasping of a set through a listing of its elements (Section 8.2.2.2) would certainly seem to suggest that. Or perhaps it is ordinality that is somehow constituted as an ordering *of a set*, one that is necessarily always in the background. Equicardinality itself may not necessarily require orderings, cognitively. Perhaps it is constructed as an individuation upon quantitative comparisons as supported by the approximate number system (which is 'continuous' and does not rely on orderings) (Feigenson, et al., 2004; Dehaene, 2011). A child may bring an (approximated) bunch of forks to the table for the guests to take, only to discover that there are not enough. And even in the purely discrete, for example, the child may grasp that many forks spread around the table are each matched with its own knife (or else some non-pair

[88] While the angle is ideally meaningless, its rotation can matter, and thus so can the action of giving it a spin. We can leave this aspect aside for the purpose of demonstration.

would 'pop out' as in some visual search tests[89]), without this perception going through any particular ordering (consciously, or even cognitively at any level) (Fig. 13). Such examples would further support the case for the cognitive double disassociation between ordinals and cardinals and its extent. There is no doubt that ordinals and cardinals (conceptually) emblematize certain operations and deductions. But if we are to take them seriously *qua objects*, then this would all seem to suggest that each is its own object-type, brought about by its own procedure array (P_{Ord} or P_{Card}). In particular, even if one of them happens to support the other in some sense, from the objectification system’s point of view, order in/dependence is *integrated* into the object-types themselves.

Just like for P_{Ord} and procedure arrays in general, P_{Card} (and its acquisition) does not simply come with its *domain*, represented in the objectification system (even if there is the domain theoretically *determined* by it). The schematic conception of what a cardinal is does not entail the holistic perception of what cardinals there are; of what *the cardinals* is, as a structure. And just like P_{Ord} in particular, P_{Card} does not settle or even represent the issue of whether there are infinite cardinals or not. Not until infinite sets are introduced and, through Cantor’s famous equicardinality abstraction, are mastered, learnt to be perceived as having or representing cardinalities.

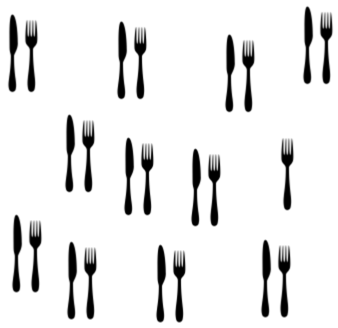

Fig. 13 Equicardinality could possibly be grasped w/o any apparent ordering involved, having a breach 'pop out' (a familiar cognitive phenomenon).

7.6 The Amalgamation

The interaction between ordinals and cardinals is substantial in general. In the *finite*, ordinals and cardinals share particularly a lot – mathematically, to the point of seeming as (two aspects of) one and the same: the natural numbers, ℕ. Let us now cash this out through the objectification framework. This section focuses on the formal, mathematical aspects of the situation, while the next one then adopts a developmental view.

---

[89] "Visual ‘pop-out’ refers to the phenomenon in which a unique visual target (e.g. a feature singleton) can be rapidly detected among a set of homogeneous distractors". (Hsieh, et al., 2011)

Speaking generally (infinity included):

The ordinals are linearly ordered. For the sake of illustration, assume AC (which holds in the finite), so that the cardinals can be linearly ordered too. Accepting that these are essentially different types, we can start by envisioning each to be on its own dimension (before the inter-regularities enter the picture). Compare, then, to a physical object, a bead, moved along a single axis, as discussed in Section 6.3. To be sure, there are drastic differences (of all kinds) between this type and the ordinals and cardinals. Besides the obvious differences regarding abstractness, a central one is in what the ontology is taken to be: The physical object, when moved, was considered to be the same object; meanwhile, ordinals along the ordinal "line" are together considered to be *different* objects, each one timeless and fully determined (and likewise for the cardinals – for the rest of this paragraph). Moving the physical object changed the parameter of its position. The nearest analogy we have for ordinals involves operations of two of them, producing a third one. Yet, an ordinal can be thought as having a one-dimensional parameter (of its location) just like the physical object's x-axis. For the sake of the illustration that follows, this dimension analogy is suitable (despite this important difference and others). The analysis concerns the *interaction* between the dimensions, not the details within each one.

We have the two procedure arrays that support the ordinal and cardinal object-types, each with its own "dimension" then. The central question arises: Could the two types be reasonably *amalgamated*? (And to stress again, an amalgamation's dimensions need be nothing alike). Mathematically, there is definitely substantial *interaction* between these two dimensions, to the point of them customarily being reduced to one shared dimension in the finite, or the ordinal dimension in the infinite (misleadingly, as discussed). The situation, in this respect, could be compared (illustratively) to the amalgamation example of the bead on a wire. In a similar manner, the level (and finer structure) of the regularity that there is to the interaction determines the amalgamation, its high-levelness, complexity, learnability, and its very possibility.

7.6.1 Amalgamation in the finite

Let us begin with the case of the finite, putting aside all facets of the general interaction. In this *restricted context*, the interaction presents an extra regularity: Taking *some* order of an order-type's cardinality always produces the same order-*type* – that of the original (as here there are no others). We can move back and forth between an ordinal and its matching cardinal – the objects are matched 1-1. And we can switch between ordinal arithmetic and cardinal arithmetic – these are coordinated accordingly. In the bead-on-a-wire illustrations, this is akin to the simple diagonal wire (Fig. 14). And the amalgamation, as was there, is comparatively simple (in this respect). Its coherentation needs to implement this coordination. For example, adding as *cardinals* the 7 and 5 objects (objects which have both ordinality and cardinality components) should now also update the *ordinal* component to 12 (with the cardinal one already managed as before). Then, 12 should (for example) be the expected result of counting-qua-ordering. And more richly, Steiner's analysis of the commutativity of (finite) ordinal multiplication (as in Section 7.2) could be substantiated. Counting up to four, three times in a row can be reframed as counting up to three, four times in a row through switching to cardinality and back. With this coordination, the two components of the amalgamation effectively act as one – and can thus be reduced to one (in terms of the implementation). The amalgamation goes from richer to *higher-level* (p. 158).

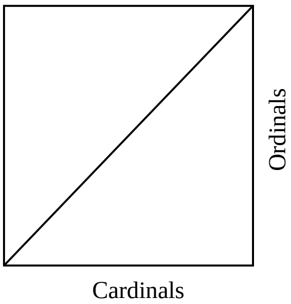

Fig. 14 The regular interaction between ordinals and cardinals in the finite (analogously to Fig. 6).

This amalgamation of the ordinals and the cardinals is precisely the object-type[90] that we come to know as natural numbers ($T_{\mathbb{N}}$). The simplicity of this coordination, in essence, is what is lost when we transcend beyond the finite, to the general case, to which we shall turn next. Just like for T_{Ord} and T_{Card}, however, $T_{\mathbb{N}}$ does *not* by itself, by default, involve any grasp of ℕ itself, the structure of *the* natural numbers. Let us return to this point and compare it with the general case after attending to the latter. For now, we already have here is a blueprint of the natural numbers that differs from the familiar ones (section 7.3). I shall pose Samuels, et al.'s (2017) question, then:

[90] As in fact the *type of object*, metaphysically, but there's no need to commit to that here.

> What are the natural numbers?
>
> Possible Answers: a) Cardinal numbers; b) ordinal numbers, c) something else.
>
> – E.g. some generic or sui generis objects, or perhaps places in the structure shared by the finite cardinals, the finite ordinals, and countless other systems

My basic answer, the fundamental representational construct of the framework I have developed, is of course c) an amalgamation of (a) and (b). Constructed upon lower-level object-types, natural numbers are thus definitely not sui generis objects. Nor is the amalgamation simply the intersection of ordinals and cardinals as types of objects in the (mathematical) world. In my account of both objectification in general (Chapter 6) and particularly with regard to these mathematical object-types (this chapter), I have stressed, and continue to stress here, an issue that is concealed in the categorization view, that merely divides the world's objects according to a set-containment hierarchy. It is the *interaction* between the intersected types of objects, intrinsic to objects of the type of the intersection, that must be integrated into the representation of the latter, as a cognitive object-type. It is not just the places they share but the *interaction* between ordinals and cardinals that is intrinsic to natural numbers, and this interaction must be integrated into their representation. Simply type-shifting between the types (another, linguistics-based possibility considered by Snyder, et al. (2018 pp. 93-95)) is not enough. We do not merely increase the length of a sequence by one, then view its positions as a set, and see that their cardinality has *happened* to increase by one. This is intrinsic to numbers as such, and it is a necessary part of our full understanding of them (including making the appropriate *prediction* here). Although finite ordinals and cardinals do by construction "share" (or rather match) a place, they do not *merely* share it. This is in stark contrast with those "countless other systems", such as the Zermelo or von Neumann finite ordinals, the uttered numerals, the people on that escalator, or my tasks for the day, à la *structuralism*. These latter instantiations can be *perceived as* numbers, abstracting them from any other content. But the relationship of natural numbers to ordinals and cardinals is different, and in fact reversed (i.e., numbers can be perceived as cardinals and as ordinals). This is not to rule out the possibility for amalgamation of natural numbers as an object-type with new object-types yet (nor that there are more abstract types integrated into them to begin with). To the extent that such a further amalgamation would keep fixed the ontology of individuated objects (rather than restrict it, e.g., to the even numbers, or extend it, e.g., to the rationals), we would take that object-type to still be about *natural numbers* (in the world), of which we have learned something new. But this is just a special case within the broader formal structure, and the finer distinctions that my cognitively grounded framework makes must not be ignored.

What of the applicability of the natural numbers? For Linnebo, as numbers have also ordinal rather than only cardinal applications, "Frege's constraint does not settle which of these applications should be built into the identity conditions of numbers" (2009 p. 228). In my version, by design, the amalgamated $T_{\mathbb{N}}$ is set to uphold an extensive form of Frege's Constraint: First, the fundamental "empirical" (through the cognitive) applications of ordinals and cardinals are built into their formal characterizations as procedure arrays. The amalgamation, in turn, has *both* types of applications built into it; both are taken to be fundamental to numbers. (All other applications are then to pass through them as a meta-cognitive abstraction, but this would need to be flashed out). This extensive form of Frege's constraint (which can be argued for) does *support* the amalgamated view (certainly in contrast to seemingly having to favor either the cardinal conception or the ordinal). This fits our starting point in Section 7.2, taking seriously, *mathematically*, the differences between finite ordinals and cardinals, even though they are structurally "identical", different only in how they are applied.

7.6.2 The general amalgamation

In the extended, general case, the ontologies are no longer matched 1-1, and thus their arithmetics diverge. This situation is akin to the wire that for some sections rises strictly vertically along the (ordinal) y-axis – all points sharing the same x-position (cardinal) (Fig. 15). (Unlike this illustration, here this happens for *all* cardinals but the finite ones. The illustration is, well, for illustrative purposes only). As in that example, the vertical sections might cause the object-coherence to break – unless the irregular interaction between the dimensions is carefully managed (and in particular, learned). For doing this, for the sake of a coherent amalgamation, the two dimensions cannot effectively be compressed into one. As with the decorative inscriptions (p. 156), once the very same decoration style can be differentiated

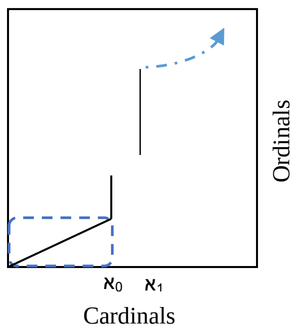

Fig. 15 (Sketchily:) The general interaction between ordinals and cardinals, embedding the finite domain and its more regular interaction (analogous to Fig. 8).

in terms of different texts, the latter are better considered to be a different object-type (or alternatively, a different component – "aspect" – of the amalgamated type, like physical engravings). This brings about an amalgamation that is richer but not as *high leveled*, which renders it *less useful* than in the finite (Section 7.2). The richness, recall, is not an advantage in itself but quite the opposite, an overload on the information management; better just keep the types separate.

Relaxing the cardinal linearity assumption (AC) – as is conceptually appropriate, as it is not essentially part of cardinality – just complicates the situation. A decoration style can actually have *no* textual content. Likewise, with AC relaxed, a cardinal can still match more than one ordinal – or it can now match *none*. This obliterates the "switch from a cardinal to some ordinal representative" ability completely. Not being even in this "one can in a sense be taken as a representative of the other" kind of relationship forces the two types even further apart (in terms of connections that are easy for the cognitive system to automate into "identity").

As mentioned previously, $T_{\mathbb{N}}$ does not by itself involve any grasp of ℕ, as a whole domain, just as with T_{Ord} and T_{Card}. It does, however, not only determine it[91], but *allows* for its perception, in a manner that T_{Ord} and T_{Card} do not seem to afford. The generation of (representations of) objects through $+1_{\mathbb{N}}$ is coordinated with all other procedures, of both P_{Ord} and P_{Card}. But the tight coordination between the latter seems to *preclude* any object not coordinated with the *generation* from through $+1_{\mathbb{N}}$. That is, $+1_{\mathbb{N}}$ is the only way to keep the 1-1 matching between ordinals and cardinals. There could be no new procedures like the one that generates ω. This forces a restriction of the actual domain (i.e., of all potential, coherent objects of $T_{\mathbb{N}}$) to the sub-domain *generated* by +1 (i.e., potential finite numbers). This does not preclude non-standard models of ℕ. And it certainly merits a much deeper analysis[92]. However, it does point to an essential difference between the structure grasp of the natural numbers versus those of the ordinals and of the cardinals (which seem to fail or offer many extra challenges regardless of compactness). Structures aside, let us return to the object-type itself.

---

[91] or "should" determine it. The relativity of infinity itself, logically driven but phenomenologically implausible, and where it fits within this framework, is beyond the point of this chapter.
[92] Which is also outside the scope of this chapter.

All of the challenging, more partial regularity that there is to the interaction between the object-types beyond the finite could still be managed mathematically, consciously, as mathematicians in fact do (even if they do not "amalgamate" because it is no longer of value). In particular, to each cardinal there should be associated the whole (possibly empty) *set* of equicardinal ordinals. This is the way things actually are between the cardinals and their possible well-ordered types. It is the "cheat" of directly associating a cardinal with its only representative ordinal that is no longer available beyond the finite. This means that what is now required is a richer set-theoretic *apparatus*. Even though some may already be required merely for the ordinality and cardinality types themselves, further set-theoretic depth does not come freely to the objectification system (which is *not* a conscious working mathematician). It is not idealized, but a real cognitive system (or a formal specification thereof). In general, each extra step needs to be acquired and managed as such; In our case, this extra layer of set hierarchy, the coordination between a cardinal and a *set* of ordinals rather than just an ordinal.

Let us compare this amalgamation to similar interactions between types of objects that are to be found in mathematics. Consider *groups* and *topological spaces*. A set can be a space endowed with both structures: It may have some operation G, a collection of triplets for which the group axioms hold, making it a group. It may also have a topology τ, a collection of subsets, for which the open sets axioms hold, thus making it a topological space too. It functions as both types of structures (*has* both structures to it), and both of these "aspects" of it can be explored mathematically. These may be independent, exhibiting little or no regularity between them. But in special cases (on which mathematicians tend to focus), they may interact, in a regular way: G may be "in accord with" τ; be continuous (itself as well as the inverse $^{-1}$ defined through it) with respect to τ. This makes the space into a *topological group*, a type of structure that is not only central enough, widely applicable, but also rich enough in theory, to have become a vast field within mathematics. The value of this "amalgamation" is due directly to the *interaction* between the aspects (or the existing *general* theories of each one would suffice). But of course, we do not just forget about each one as a separate object-type. There is independent value to each one, not only because i) each may be theoretically developed on its own (where keeping in mind the other too would just be a distraction), but more importantly because ii) each one may certainly appear, be applicable, separately (or sometimes the two together but not coordinated, without the aforementioned regularity between them). In comparison with ℕ, the difference (in this respect) is mainly one of *degree*. Within the general scope of applicability, the subdomain in which ordinals and cardinals are tightly coordinated is much more central, immensely more applicable – and basic. From this section's formal (objectification framework) perspective, this is not an essential difference – but it makes all the difference in the world when it comes to the developmental tracks of how and when we come to acquire such object-types.

7.7 Mastering ℕumbers

I have portrayed the different object-types and the relations between them (in the finite and the general contexts). What we further need to understand is the developmental track the child goes through, the dynamics of her route through this space of procedure arrays: What arrays she picks up on directly, what amalgamations she constructs to reach new arrays, which arrays she keeps and which she discards. I have established the general principles that guide such developmental dynamics in Section 6.5. But the full, actual picture cannot be grasped through such rather a priori analysis alone. The facts of the matter depend both on the particulars of the cognitive makeup and on the environment through which the child is exposed to the relevant regularities. The existing empirical work might not be of enough use here, as described by the perspective of a philosophically and mathematically informed source:

> Developmental psychologists have tended not to pay much attention to the cardinal/ordinal distinction!... some of the most prominent evidence for the possession of number concepts is neutral with respect to interpretations of those concepts as representing cardinalities or ordinals… the fundamentality of cardinality seems to be an assumption adopted by psychologists, *qua* theoreticians, in framing their research and interpreting their data (Samuels, et al., 2017).

(For a psychologically informed source, see (Rips, 2015)). Fortunately, we have enough to go on for framing the issue, and hopefully suggesting a reasonably likely picture, particularly (to stress again) for the *relations* between the types:

In the (cognitive) beginning there is nothing: no numbers, no cardinals, no ordinals (and quite possibly no abstract object-types of any such kind). For cardinals, or numbers, as I have described at some length for ordinals, the system can come to detect the relevant regularities as being about the general type *object* – whatever they may be.[93] However, objectifying successfully is (as always) a search problem within a vast space of possibilities. The system needs to be exposed to the relevant regularities. And furthermore, the system (though not necessarily the person, consciously) needs to attend to the regularities – *with respect to the relevant set of procedures*.[94] Homing in on these types in question, as in general, can:

- possibly never succeed.

- "succeed" trivially, because the types are innate, already there to begin with or in a fixed development course. But for the full type ℕ in all its glory, including the adult, *more* (if not perfectly) conscious understanding of how to apply them and reason about them, this seems rather unlikely.

- succeed, unsupervised, because the types are simple enough and/or the system is attuned enough to noticing the regularities and capturing the types (like the baby's propensity to focus on faces – even before it has probably learned them as the object-type that kids come to know).

- succeed through supervised learning. This is the standard course by which we usually come to learn the types. They are not introduced and cognitively mastered through any explicit definitions (e.g., in second-order logic?); they never are fully specified consciously (i.e., with ontological rigor). Rather, they are taught by exemplification, through perceptions and manipulation. When singled out (i.e., by a teacher), they can more readily be identified by the system as coherent object-types; the system is 'encouraged', brought, to focus on particular procedures and arrays. Merely teaching together different embodiments of addition (apples, cows, etc.…) and using the same words for designating the actions and results, for example, can turn the spotlight away from the particulars of each embodiment to the (meta-cognitive) regularities from which these are abstracted away.

---

[93] I have posited that the *objects* are what the objectification could come to accept as such, and that the regularities are accordingly in fact cognitive, though this is not essential here.

[94] Abstract (meta-cognitive) types may be harder to home in on, as they have as prerequisites familiarity with other less-abstract procedure arrays.

The actual details, for these discussed types, fall somewhere in between all of these. For example, "matching things" may perhaps be a more natural activity than figuring out the particular order (perceived as a completed ordering, not just the ability to, say, order people by height by simply comparing two each time). So, it may be that people could easily, always, pick up cardinality, whereas learning to recite things (i.e., in order) requires more guidance. There is also room for variance among people here. Variance can result not only from differences in the systems' general "ability" to perform its task, and the general environment (which may include better teachers), but also due to differences in inquisitiveness; in the propensity to explore, strive to expose one's system to certain types of procedures. At its most extreme, this surely had a part in the first discoveries of such types – before there was anyone else from whom to learn of them. The objectification system always plays a similar role (or at least there is no reason for us to assume otherwise). Ultimately, for the system, these are all matters of degree.

Let us again turn to our main focus, on the relations among the three types. *Formally* (i.e., as to the mathematical object-types), I have made the following points:

1) Ordinals and cardinals are essentially distinct object-types, even in the finite.
2) Their amalgamation happens to be simpler (and thus easier to master) and higher-level (and hence useful) in the finite.

Developmentally, I have argued, in general, for a fundamentally *bottom-up* picture (p. 168). Consequently, I now suggest that the cognitive objectification, the learning, of these particular object-types, undergoes a similarly structured course ("development recapitulates… ontology"). At least, this is the rough sketch, tracing the fundamental drive (where the finer details could be explored from this point on).

In the environment, in the background, there are the parents and teachers, who bias the regularities and direct the children towards mastering numbers. Their strategy for achieving this balances between how children actually best come to master numbers – and how the adults perceive them. As the ordinal-cardinal distinctions on which we have been focused are not consciously well-perceived by adults, their teaching, too, is *biased* away from the differences between the types, into teaching directly about the amalgamated numbers, both aspects of their application included. The bias still leaves room for partly separating the two "aspects" if this is how children best come to master numbers. This indeed is part of my substantial developmental contention here.

Consider first the possibility that (at least rudimentary parts of) the procedure arrays P_{Ord} and P_{Card} are at some early point acquired, individuated, as independent ones (even if sharing many procedures or sub-computations between them). They are not about *finite* ordinals or cardinals. Even though all of the examples encountered (for either) are in fact finite, the issue of finitude plays no part and simply remains unspecified. The (Kripkean) child (p. 185) could come to see that cardinal multiplication (axes switching) is commutative while not seeing why this should be the case for ordinals because only the former is inherently permutation-independent. To be sure, the commutativity in ordinal arithmetic *is* a regularity in the finite, and so potentially could have been assimilated directly into P_{Ord}. But commutativity is not trivial to notice (for multiplication more than for addition); at least it is not trivial upon first acquaintance with this notion, before one is already familiar with it and poised to look for it. Most importantly, seeing *that* cardinal multiplication is commutative (in that sense) is not merely becoming accustomed to it, assimilating it into the type; it is seeing *why*, seeing a reason *for* it (which *can* in turn build on other facts one did simply get used to). With enough exposure, the child perhaps *could* come to assimilate (through weaker, statistical care) ordinal arithmetic as commutative. But the easier route, of seeing a *reason* for it, passes through cardinal arithmetic. As the amalgamation that supports it is, too, something that the child will soon be brought to master, she never needs to absorb commutativity into P_{Ord} itself, directly (at least as a reasonable developmental route, ruling out e.g., 'unnatural' teaching approaches).

Of course, P_{Ord} and P_{Card} never come to be fully mastered (by the child), in particular as their different nature in the infinite is far from even beginning to be explored. However, this goes for the natural numbers, too, for which even a professional number theorist's mastery, indeed her "perception" of them, would continue to develop throughout her lifespan. It is a matter of degree. The fundamentals for grasping numbers are considered to be in place by some early childhood stage, and so the fundamentals of cardinality and ordinality probably are there too.

Where the extra regularity of the finite takes effect is, as I have analyzed at length, in coming to amalgamate the types. The particular mapping between the types, which is possible within this limited context, stands at the core of the higher-level amalgamation that the child comes to know as numbers. Being taught how to count instills this mapping by stressing *abstraction* of the reciting result not only from the content of what is being counted (whether apples or cows), but from the order by which it is done. Order-irrelevance is an in-house P_{Card} matter; *turning* ordering irrelevant (and back again) is a manipulation the child must swallow for amalgamating P_{Ord} and P_{Card} (in the finite) into $P_{\mathbb{N}}$. This back-and-forth manipulation and its (perfect) regularity should become so ingrained as to become *automated*, no longer exerting a cost upon conscious processing.

The central educational (if implicit) tool for this is the *linguistic designation* (p. 163). Relating the *seventh* ordinal (or place) to the amount of *seven* (etc.) not only directs attention towards the object mapping thereby facilitating its detection, but it effectively *enforces* it, *systematically*. Linguistic convergence not only relates but binds together ordinals and cardinals – even if they were consciously strictly differentiated to begin with! The finer difference in linguistic expression still reflects ordinality and cardinality as *aspects* of the role of the object in the cognitive system, the full-blown concept of number. But its automatized systematicity (once mastered) instills the mapping.

A linguistic account of ordinals versus cardinals to complement this picture would be greatly desired here, particularly with regard to how the system's acquired object-types are reflected linguistically. Even a brief pedagogic account of the ways by which the child is brought to master the amalgamation is beyond our scope as well. What is most important for us is to see the teaching of numbers in these general terms, as an extra layer of optimization that occurs within the objectification system and rests *upon* the *formal* structure of the types at hand and the relations between them.

Intertwined with the pedagogic details are, as aforementioned, the developmental ones. An account of them, accordingly, is beyond our scope too. But let us dwell on them a moment to see how the framework and my own account relate to that literature. The particulars of the cognitive makeup constrain and guide the learning path and can determine what comes first and how things come to be. Alternatively to my first suggestion above, it is reasonable to suspect that at first only P_{Ord} comes to be acquired, followed directly by the construction of the full amalgamation, to $P_{\mathbb{N}}$ itself. Even if P_{Ord} and P_{Card} can be formally differentiated, perhaps the main or only feasible route to understanding equicardinality is via counting, through grasping particular orderings and then symmetrizing them. For another alternative yet, Berteletti et al. (2012) offer results that "suggest that the child's conception of numerical order is generalized to non-numerical sequences and that the concept of linearity is acquired in the numerical domain first and progressively extended to all ordinal sequences". However, "It's currently unclear whether the extant evidence from developmental psychology resolves issues regarding the relative developmental priority of cardinal and ordinal concepts" (Samuels, et al., 2017). It may be that any such schematic suggestion is ultimately simplistic, with the actual details more complex and gradual. Each simpler types' coherentation and their amalgamation into ℕ may all be intertwined. What is empirically evident is *the struggle to master the amalgamation*. Arriving at the fundamental Cardinal Principle comes gradually:

> Carey showed that different 3-, 4-, and 5-year-old children can be classified as 'one-, ' 'two-,' or 'three-knowers,' or as 'counting-principle-knowers.' The one-knowers do well when asked to provide one object but not when asked the other numbers, while the two- and three-knowers can respectively provide up to two and three objects successfully. 'Counting-principle-knowers' count quantities above three or four as well. (Bryant, et al., 2010 p. 554)

From there on, the child knows that the last numeral expressed answers the question of "how many?", as a general rule. This stage, too, although central, is only a partial one, after which the struggle still continues. For example, as Davidson et al. (2012) show, "these children often do not know (1) which of two numbers in their count list denotes a greater quantity, and (2) that the difference between successive numbers in their count list is 1". Bryant et al. (2010) also offer various examples of the complexities and stages in coming to master numbers, integrated with the complexities of exploring it empirically. Ultimately: "Virtually all the … evidence on children making number comparisons supports the idea that preschool children do not yet have a well-developed understanding of number." (p. 556)

This perspective, with its attention to the educational challenges and the possible failures, widens the scope not only of what cannot be innate, but what cannot be picked up easily (without guidance) by a general objectification system (and perhaps could not at all): "We cannot take these and other basic logical principles for granted. They are a source of genuine difficulty for children, and the idea that they come as an innate and universal gift is misguided and actually harmful, for it distracts us from giving help to children where they need it". (Bryant, et al. p. 569)

Meanwhile, at the other end of the spectrum, empirical explorations of "the number sense" and other natural seemingly mathematical abilities widen the scope of what is shared by all (some animals included). Such abilities may be supported innately, as is the common working assumption in the field. A different possibility, however, is that they are simple and evident enough throughout life (or important enough to bias perception towards them) that a general objectification system, or some non-object-centered one, would necessarily come to pick up on them.

Relating my theoretical framework to the actual details and many empirical findings, done properly, would require much more than I can offer here. However, framing the topic as I have done places emphasis on the issue of ordinals versus cardinals, an issue that has largely been neglected throughout the vast relevant literature. The previous standard developmental psychology source (selected as such) demonstrates this well:

> Knowing that a set of 20 items has the same number of items in it as another set of 20 items and that a set of 20 items is more numerous than one of 15 items is relational knowledge. "Cardinality" is the correct term for the first of these relations and "ordinality" for the second. Cardinality means that a set with a particular number of items contains the same number of items as any other set with that number. Ordinality means that numbers come in an ordered series: 3 is a larger number than 2 and 2 than 1, and it follows that 3 is also larger than 1. (Bryant, et al., 2010 p. 550)

In standard *mathematical* terminology, "a set of 20 items is more numerous than one of 15 items" is a purely cardinal relation. It can be deduced by surveying all possible mappings (in no particular order), though it happens to have an effective procedure (counting) to determine it (in the finite) that *passes through ordinality* (i.e., from "shorter than" to "fewer than"), as has been previously discussed. The $15_{\text{Card}} <_{\text{Card}} 20_{\text{Card}}$ relation itself *is* an ordering (of cardinals), in which equality (equicardinality), too, is part of the order. (Would $\leq_{\text{Card}}$ be a cardinality relation? Ordinality? Both? Neither?) The transitivity of "larger than", for its example, is purely *about* cardinality (as opposed to "follows"). It is only the *form* of the principle, its (general) entailment structure in which "it follows that 3 is also larger than 1", that goes through orderings (of propositions).

What could explain such "alternative" terminology? A developmentally centered view may very reasonably put the concrete before the abstract. Whatever other generalizations of orderings interest mathematicians, perhaps a child's basic familiarity is with orderings is always through numbers? In a sense, this is right: mastery of numbers is where the concept of *order* may first be made explicit (if not in its abstract form). But as detailed in Section 7.4, orderings are quite fundamental, cognitively, and accordingly, could perhaps be mastered to some degree before or independent of issues of cardinality. To be explored as such, however, they must be conceptually differentiated from cardinality from the start. Otherwise a child's basic familiarity with orderings *may* seem to always be through numbers simply because numbers (the amalgamation) are taken to designate places in an order too. The only paper surveyed in (Bryant, et al., 2010) that mentions ordinality, for example, is (Brannon, 2002). It explores infants' perception of order – but still, an order of quantities (differentiating whether the sequentially presented cards show an increasing or decreasing number of objects). That too is of course integral to the full mastery of numbers; but to understand these integration issues, perception of on order should also be explored independently, if possible.

Even this general, inadequate attention to ordinality for its own sake, is as part of a call for a finer attention to our topic than the standard in the field:[95]

> The distinctions that we have drawn between quantities, relations, and numbers are not new. Piaget stressed the importance of children learning not just about counting but also about numerical relations, and in particular about cardinal and ordinal number. More recently, Thompson (1993) argued that people must be able to connect numbers, quantities, and relations in order to reason mathematically. Nevertheless, many recent psychological studies of children's mathematical learning have bypassed

[95] For a review of studies that do pay attention to ordinality and its in/dependence from numbers, see (Sury, et al., 2012).

Piaget's and Thompson's injunctions about quantitative relations. (Bryant, et al., 2010 p. 550)

In the end, the child does come to master the amalgamation as such. She becomes accustomed to seeing both ordinals and cardinals instead as higher-level, full-blown numbers. She becomes convinced that their arithmetic is commutative through their cardinal aspect (and with time gets used to the fact), never having to wonder directly about *ordinal* commutativity, as Kripke did. Even if at some stage P_{Ord} and P_{Card} were singled out as procedure arrays that support distinct types, they come to be "thrown out" as such, rendered by themselves useless, as economy dictates (Section 6.5).

The amalgamation, and its successful mastery, are of course contingent on the central regularity, which holds only in this restricted context of the finite. The *general* non-necessity of this regularity (for orderings and amounts) is never made explicit. In every demonstration or exercise, the regularity is always there in the background, like shadows that appear (when they do) almost always below physical objects. The system can pick up on the regularity; the child comes to rely on it, automatically, like judging height versus depth of a ball (Section 6.4). To that extent, an "environmental", non-necessary fact sneaks into the very constitution of mathematical ontology. This was probably the main force behind the original discovery(s) of the natural numbers to begin with. As mentioned previously, though, the regularity is also forced upon the child through linguistic convergence into numerals. In this respect, as the regularity in a sense *defines finitude* (as the domain within which it holds), this is another central component of what the child is *implicitly* taught, over and above ordinality and cardinality and their amalgamation: finitude.

Familiarity and mastery of the amalgamation that produces the natural numbers need not yet entail a grasp of $\mathbb{N}$, as the completed *domain* of the finite (to follow up on Sections 6.7.2, 7.4, and 7.5). As suggested, this is likely the result of further meta-cognitive reflection upon the procedure array $P_{\mathbb{N}}$ itself. One needs to see that all that $P_{\mathbb{N}}$ can *generate* is actually all that $P_{\mathbb{N}}$ can *operate on coherently*. But this is not part of what the objectification system itself supports, and not part of my account here.

Many years later, that child hopefully comes to take her first course in set theory. There, the simplifying regularity of the finite, which was linguistically disguised as object-identification, is lifted, and the full workings are made explicit. Ordinals and cardinals are properly individuated, and their general interaction begins to be explored, with the necessary set-theoretic apparatus being acquired (as in Section 7.6.2). For the objectification system, P_{Ord} and P_{Card} were already there in terms of many of their procedures and their relevant regularities. However, they come to be re-designated as (object-coherent) procedure arrays, on their own (aided by the system's familiarity with them). They are revived as independently (consciously) perceivable, just like the infant's mother's voice and face (p. 170). Their amalgamation in the finite is left standing. But the general case amalgamation, not higher level (and thus useful) but just richer, is put aside for the sake of the simpler, more meager types. No longer having the matching regularity at one's disposal for free use, things have to be unlearned. The bonds that have been made and then hidden through automation need to be broken (which may not come easy to the student). Learning that ordinal arithmetic, on its own, is in fact non-commutative actually requires explicit demonstrations and getting used to.

This much later stage of mathematical development shall be the story that the next and last chapter explore. Let us now conclude with a philosophical overview.

7.8 Philosophical Summary

As the classical mathematical narrative would have it, numbers only had to be split into ordinals and cardinals when mathematics went beyond potential infinity, into the infinite. The basic insight I have adopted, that ordinals and cardinals are essentially, mathematically different in the finite too, unsettles this theory. With my objectification framework, I have offered a more systematic space in which to embed and understand these different types and the relations between them – and their developmental (cognitive) track.

The framework, with extensions and amalgamations at the center of its dynamics, imposes a bottom-up view. What matters is the richness of the procedures and the level of coordination between them – not which determined ontology happens to contain which or is generally larger. That ℕ is "contained" set-wise in both the ordinals and the cardinals (when we lose sight of their difference in the finite) does not mean that ℕ *precedes them foundationally*. And it does *not* mean that ℕ is simply "simpler", sense unspecified. The developmental route should be expected, as a general rule, to follow in the path of these foundations, to recapitulate the formal structure (though this can be overturned by tailored learning supervision, innate propensities, etc.). Bringing more procedures into richer and higher-level arrays, orchestrating them into object-coherence, requires more – more effort, examples, regularity detections, time, and so on. This is why amalgamations are a *directed* phase driving (in part) the objectification dynamics. This is the fundamental drive.

Under this view, the causal and, in fact, foundational direction is reversed, the classical narrative overturned. It is not that going to the infinite forced us to break apart ordinals from cardinals. They are distinct to begin with. Instead, they were set up for the cognitive system to bring them together through the restricted examples to which we attend and the regular interaction they exhibit (and enforced through unifying language), thereby in effect restricting to the finite. The system by its drive not only pursues this (valuable) merger but renders the original distinct notions and interactions between them implicit and in fact subconscious, since this is possible (and in that case, valuable, too). This possibility, however, is merely a statistical and contingent possibility, not a mathematical, absolute one. The ability to “identify” finite ordinals with cardinals is akin to some Euclidean geometry “fact” for which its reasoning inconspicuously depends on the drawn triangle happening to be acute-angled.

What can we learn from this in the general spirit of the project? This case study deepens the cognitive critique of a priori, non-cognitive approaches to mathematics. Specifically, where it comes to distinguishing – *within mathematics* – the accidental from the truly mathematical. Foundation-wise, the case study argues the importance of keeping to the cognitive processes themselves as underlying ontology – at the very basic level (in my terms, aspiring for *ontological* rather than just *logical* rigor): A logical formalization (e.g., into Peano Arithmetic) may (misleadingly) capture higher-level aspects of the actual ontology, which, although rooted in the cognitive processes, are *partial* to the ontology's full mathematical contents (which does not include just the top-most, conscious level).

8 The Universal Set in Naïve Set Theory

The principle of unrestricted comprehension is the main intuition, concerning what sets there are, that can arise from a mastery of naïve set theory. Yet something must be going wrong there, as it is known to lead to contradictions. I now present the case as a conspicuous example of how the mathematicians' unawareness of their own internal processes may lead to amazingly *robust* – yet demonstrably false – *fundamental* intuitions that govern the perceived ontology. My account, following some background, takes place in two phases.

For the first case study, Chapter 6 presented an abstract framework for the representation of mathematical objects, which was then applied in Chapter 7 to ordinals and cardinals. The first phase now continues this spirit, sketching a larger-scale representational account of naïve set theory, in which the non-formal notion of a set suffices. This includes both pre-existing relevant knowledge and the object-types that form as the student comes to master the topic. In this light, the discussed intuitional failure appears unsurprising. Easy conclusions in favor of a general skepticism towards such intuitions aside, this sketch serves as a basis for the next speculative, but *constructive*, phase.

I present an analogy with *visual* perception. This topic has been massively explored by the cognitive sciences. It offers a wide variety of empirical findings and illusions that reveal surprising limitations of our perception. These results stand in stark opposition to our experience of that perception, which lacks awareness of the underlying mediating cognitive workings with their riches and limitations. This contrast makes up for a 'grand illusion' of a detailed perception of the whole visual scene, which I term *omniperception*. Clarifying this phenomenon in sufficiently abstract (non-visual) cognitive terms allows for considering the possibility of an analogous phenomenon occurring in the "perception" of sets. Such set-theoretic omniperception would entail that the universe of sets forms a set. This account thus locates the source of 'illusory' intuition for that universal set (and through this, for the principle of comprehension).

8.1 The Background Story

8.1.1 Ontology Introduction

Our story is about a typical first course in so-called *naïve* set theory, and the typical *competent* student, coming to master it successfully (learning difficulties aside). It is "naïve" in the following way: Throughout the course, its subject matter is never formally introduced (not the way vector spaces or graphs or the continuum would be, in their own basic courses). No axiomatic foundation nor model are given; no definition of what sets are, what sets there are, or what sets there are not or could not be are provided. Even the axiom of extensionality, of no two sets sharing the very same elements, may be picked up implicitly. Pedagogically, this is probably as it should be; these theoretical issues do not stand in the students' path to mastering the material of the course. "Naïve" mathematics can go a long way, as it certainly has throughout the student's personal educational history as well as throughout mathematics' actual history, when being formalized only in retrospect.

A pre-theoretical conception of *a set* – of what a set is – is appealed to. Though not explicitly relevant to the course, it is based in the prior implicit acquaintance with great many real-life sets (and related constructions and concepts) that inhabit the students' mental realm as humans: One's immediate family, the Republican Party, soccer league (as a set of sets), etc. Beyond these, the mathematical set ontology unfolds as the course progresses, introducing more and more *concrete* mathematical sets (i.e., particular, attended to, individuated in working memory); various types of constructions of sets from other sets (constructions that are then generalized); and related abstract concepts. Some of these may be familiar, some are new. The introduction of concrete sets is intertwined with the introduction of related concepts (i.e. cardinality) – and proofs. Proofs are about concrete such structures (the literal diagonalization of $|\aleph_0| < |c|$), as well as about the concepts in the abstract ($|A| < |2^A|$). Importantly, proofs may rightly be considered an integral part of the acquaintance with the ontology; an integral part of its 'perception'. For example, intuiting cardinality as a notion of "size" depends inherently on the Cantor–Schröder–Bernstein theorem ($|A| \leq |B|$ and $|B| \leq |A| \Rightarrow |A| = |B|$). However, here I shall focus on the ontology directly (rather than the course's much-richer contents as a whole). In all of this, while more and more sets are introduced ad hoc, the accompanying issue of *what sets there are* is (typically) never even raised.

8.1.2 Comprehension

Historically, set theory was not a distinguished mathematical field, concerning a single notion:

> [N]aïvely sets present themselves in a number of distinct contexts (finite collections of concrete objects before us; sets of natural numbers satisfying more or less explicit conditions; sets of points in geometry). One may therefore doubt whether any definite general notion (of set) is involved here; it looked more like a mixture of notions. As a matter of historical fact this was the common feeling among Cantor's contemporaries. (Kreisel, 1969 p. 93)

In the context of the modern, more abstract tradition in which the course is taught, the situation is different. Having familiarity with quite a few concrete sets (not to mention the many from outside the course) as well as some construction rules, it is but natural to seek the unifying notion underlying them all, having them as its scope, especially as one – that of "set" – was being alluded to constantly from the very start. So, following a successful mastery of the course (and quite possibly before that, too), the student is confronted with the question of a unifying rule that governs set existence, upon which the principle of ***comprehension*** is then suggested:

For any property p(x), there is the set {x | p(x)} (of all the objects of which p holds).

The principle may seem not only perfectly plausible, but a contender for that rare status of a *self-evident truth*, and possibly the precise formalization of the single construction that has implicitly been used to introduce the various concrete sets as well as the (logic-induced) set operations all along. It produces any set that is required in the mathematical practice into which the students are taking their first steps. The reason for this is (*seemingly*) simple: Any concrete set that mathematicians would like to exist – any set they may ever come across or find necessary for the mathematics they are doing – necessarily has a way to be referred to; they would talk about it. In particular, they could determine it linguistically through the characteristic of its elements, a predicate – and thus have a matching property to apply comprehension to. (The axiom of choice famously loosens this requirement). Language does all of the work (as with the logical foundations generally, Section 1.1), and set theory hinges on that.

The nature of this principle *as a constitutive one*, as a main commitment to set existence that the practice relies on (bar choice), a principle which has allegedly (in retrospect) been used to introduce ontology, is delicate. The principle would not be able to *directly* insert into the ontology all of the sets the mathematician would like to exist, given reasonable (language-grounded) restrictions on the ontology of properties: The denumerability of the language's expressive power would prevent it from naming most of the particular subsets of $\mathbb{N}$, for example. However, as comprehension does give "all of $\mathbb{N}$'s subsets", i.e. $\wp(\mathbb{N})$ – it does insert them into the ontology, albeit *indirectly*. This is a metaphysical leap, from a human-relatable foundational basis that, accordingly, is not *object-complete* (p. 18).

8.1.3 The Universal Set

In the name of unificatory elegance, comprehension builds upon properties of sets and of objects in general – *any* (seemingly well-defined or mathematically-coherent) property. In particular, one can use comprehension to immediately produce the following:

1. U, the set of all sets, using the simplest of properties, which simply holds always, for any set.
2. $\wp(\text{U})$, the set of all subsets (*whatsoever*), using the also seemingly simple property of being a subset of an already given set.

U, the *universal set*, would obviously have to be the largest set. On the other hand, according to Cantor's theorem, $|\wp(\text{U})| > |\text{U}|$; its power set is strictly larger. Accepting comprehension with respect to two of the simplest logical properties, we are faced with Cantor's paradox.[96] The proof goes only through the most basic of notions and proof principles (and in particular, contra a common misconception, is perfectly *constructive*). Thus, it seems that something has gone astray regarding the ontology itself; that the universal set (when accompanied by the power-set construction) may not be a legitimate object of set theory.

---

[96] From this, Russell's famous derivationally shorter paradox can be (and was) extricated (Ferreirós, 2007 p. 286), but it is more cumbersome in terms of the point of view I am taking.

Presented this way, U is but a very specific application of the principle of comprehension. But as it uses the simplest of properties, it is actually paradigmatic. (The other known set-theoretic paradoxes seem to be of a related nature). Indeed, it is standard practice to get new sets from existing ones by filtering in only those to which some predicate applies – both in naïve set theory implicitly (e.g. "the even numbers") and in most axiomatic versions explicitly, through the scheme of *separation*[97]. So, unless separability itself is contentious, then U, ontological commitment enough, is all the illusory intuition we need to explain.

It should be noted, from the point of view that is developed here, that separation should indeed be approached with caution (rather than be taken as an obviously safe commitment), for it goes against the following fact: A substructure may be more complicated (for various senses of "complicated"), may be embedded more complicatedly, than the structure in which it is embedded.[98] But whatever stance one takes on this, even just U is already too much anyway.

A further caveat would concern the power-set construction itself, whether granting comprehension or not. From a cognitive rather than a logical point of view, it becomes anything *but* obvious (and must be accounted for; see Section 8.2.2.5). In particular, it is the very reason a cognitive foundational basis can no longer (i.e., for post-Cantorian mathematics) be object-complete. But as discussed previously (Section 4.3), my approach relies on the benefit of historical hindsight rather than take on revising of standardly conceived current mathematics at this stage. Thus, let us just accept the power-set construction and keep the focus on U as what needs to be explained (away). Where does the intuition for it come from?

Platonism, admitting an utter metaphysical disconnection of the objects from their representation, is well known for supporting tolerance of impredicativity: the objects are there anyway, and our definitions (or cognitive perceptions) simply manage to identify and refer to them. This works just fine for the least upper bound:

---

[97] Also known as “specification” – and as “restricted comprehension”.
[98] Accordingly, alternative set theories (such as Mac Lane set theory and constructive set theory) may indeed restrict separation (usually to formulas in which all quantifiers are bounded to sets – “Δ_0 separation”) (Holmes, 2017).

> in mathematical practice, the real numbers are regarded as a definite, completed totality independent of human constructions and definitions, not some vague collection of numbers "growing" under successive definitions. Then the l.u.b. [least upper bound] of a bounded subset S of ℝ merely serves to single out a specific element of ℝ, not to "bring it into existence". In this respect, it is like the least number operator in the set of natural numbers. (Feferman, 2005)

What of that universal set-of-all-sets, then, or Russell's infamous subset? After all, these are the very archetype of impredicativity (with U singled out through a quantification ranging over itself too, amongst all other sets). For the least upper bound, there was an operator quantified over an ordered set of reals to select the minimum (set-theoretically, through a union/intersection of Dedekind sections). Comprehension supposedly serves as the analogous operator used to produce U: It quantifies over this ultimate collection of sets and selects the set that this collection *supposedly then is*; committing to it as a set. Apparently, there is much more here than selection merely among the pre-given totality that is ℝ. The plurality of *pure sets* is assumed to be a definite, completed totality, a collection that can be quantified over. This alone poses no problem for platonism. Cognitive realism too (Section 2.6) may support such a definite totality. Even a neo-psychologist may perhaps accept a plurality that is well defined with respect to a given biological rooting. Non-realistic cognitive approaches (e.g., Longo's, Section 3.4), however, would reject such a totality altogether. The intuition for comprehension and U, first of all, builds on some form of (mathematical) realism.

Being able to quantify over a totality, however, need not automatically entail the *set-hood* of that totality (which, i.e., can be an element of other sets). Rejecting this is how the modern mathematical platonist, while giving up on comprehension and U, nonetheless keeps to set realism through the class-set distinction[99]. This is how the practice itself came to avoid the paradoxes. But what is the nature of the distinction between classes and sets? Ironically, for a field that was supposed to be an absolute, preexisting abstract bedrock for everything else (temporal series included), the accepted solution adopts a *generative*, "constructive" approach. The principle of the *limitation*

---

[99] Though there are problems with such a conception. See e.g. (Shapiro, et al., 2006) and others in (Rayo, et al., 2006).

of size rejects – methodologically – the platonic reliance on a preexisting definite domain in which objects can be collected into sets using arbitrary properties. Instead, ontology according to it is tracked (if not "built") from the ground up, through constructions (e.g., union, the power set...) that we ideally could follow in our mind's eye. The metaphysical platonist can still take these constructions to merely reflect the structure of the preexisting definite domain; causal construction directionality need not entail metaphysical directionality. But the construction is inherent to the determination of what sets are valid. (The later *iterative conception of set*, conceptually related to but not completely compatible with the limitation of size, makes this explicit).

Even when one comes to accept this alternative limitation-of-size or iterative view of the set universe, it is a theoretical acceptance. It need not be an immediate realization, *natural* in retrospect (an "a-ha!" moment), that this in fact reflects what has been going on throughout the naïve course. (Comprehension, in comparison, is much more intuitive to begin with). How is it that the student, having successfully mastered the course, can nevertheless miss out on such a crucial supposed distinction from classes, locating the *constitution* of sets elsewhere completely?

8.2 A Cognitive View of Set Theory

If a naïve, unformalized notion of sets can suffice for the general mathematician for the set theory he or she requires and can work with successfully, then its mental grounding should ultimately be accounted for. For the first phase of this case study, I now sketch such a large-scale representational account of naïve set theory. The first subsection addresses the possibility of preexisting (possibly innate) relevant knowledge. The second subsection then attends to the object-types that form as the student comes to master the topic. Finally, the concluding subsection draws an early, general skeptical conclusion towards intuitions like the one for the universal set.

8.2.1 Set Theory in General Cognition

Officially, our story is about the university set theory course. The actual cognitive story of the concept of set, however, begins long before that. This is analogous to how, recall, some fundamentals of ordinals and cardinals (Chapter 7) may have been laid out long before the child is confronted with numbers explicitly (having to perform calculations and master numbers at a preliminary *theoretical* level). Sets are a part of a mathematical toolbox, which may be useful in cognitive life, for humans at large and animals too.

This means that A) the general cognitive mechanisms may come to grasp them before the explicit pedagogic stage, and B) evolution may have come to do some learning of its own and assist through innate mechanisms. What was true for numbers (and their constituent object-types), may be more far-reaching for sets. Abstract and logically fundamental as they are, sets (if not set *theory*) and their manipulation may well have a role at the service of general cognition, not just for the explicitly mathematical (i.e., in mathematical education):

The obvious starting point in the search for sets within cognition is in the framework for the semantics of the mental world (Lucy ∈ dogs ⊂ mammals; the families in the village, etc.); sets as part of (if not constituting) natural categorization. However, less-obvious applications may also abound. Real-world objects may be conceived of as (structured) sets of smaller objects (a chair having four legs) which may, in turn, be best recognized as (structured) sets of components still (edge detection). Do sets, then, play a role in the *visual* system? Gestalt *grouping* principles, for example, seem to concern this explicitly.

Fig16 . "The Gestalt law of proximity states that objects or shapes that are close to one another appear to form groups". *(Wiki-PoG, 2018)*

Caution is required, however, in relating the mathematics that subsystems of the brain may be doing to the ultimately explicit, conscious mathematical matter. The fact that a set of neurons in a cat's or a toddler's brain may be computing some derivative, say, does not at all mean that derivatives are part of their mathematical capacities in any meaningful sense (for us). *Of course* the brain performs highly sophisticated mathematics, much of which science probably has yet to discover. Furthermore, even some cognitive capacity being necessary for some agreedly-mathematical capability – is still not sufficient for being relevant or illuminative of mathematics' cognitive foundations (breathing might be necessary too). In claiming such relevance, the burden of proof is a tall order.

I suggest that mathematics-performing brain mechanisms become more relevant the higher-level they are; the closer they are to consciousness. In particular, following the model(s) of the "global work-space" (Prakash, et al., 2008; Dehaene, et al., 2011), a mathematical capacity needs to generally be available for other modules to make use of; the usage of the evolved or acquired capacity for various including novel purposes needs to be a real possibility. Put the other way around, this rules out any mathematical capacity that some module happens to implement for its own purposes but does not offer as such, in the abstract, independently of its extra-mathematical task, to other modules. An archetypical example of a truly *mathematical* module may be the approximate number system (which is innate and shared with many other animals) (Feigenson, et al., 2004; Dehaene, 2011). It is (supposedly) used sweepingly for approximation tasks in general across various modalities. When we come to master arithmetic explicitly, the acquired capacity latches on to the system. It is a common claim that the approximate number system is even a central component in our journey to mastering arithmetic. Might explicit set theory analogously build on any such pre-set faculties?

There is a bit of modern empirical research regarding sets in cognition. To bring in one: Feigenson (2011) has explored our ability to perceive real-life sets and properties about them (e.g., "most faces are smiling") in the face of the limitations of working memory (Section 5.1). Crucially (for us here), *the "sets as objects" cornerstone of set theory gets its direct cognitive manifestation*: "each set…functions much like a single individual object for WM [working memory]". Furthermore: We have the capacity to *approximately* compare sizes ("Array A has more dots than Array B") and various statistical features ("Array A has dots that are smaller on average than Array B") for large sets (of e.g., hundreds of items). This means that, cognitively, these sets cannot be represented through the standard "mathematical" idealization of a set, in which every element is individuated; working memory cannot hold enough items. That brings about a related cognitive notion of *ensembles*, which support such operations. "The non-retention of the individual items in the array makes ensemble representations importantly different from object or set representations". That common claim that numbers are constituted upon the approximate number system, together with subitizing (through object-individuation in working memory), might thus originate in the representation of the collection itself (i.e., before happening to compare their numerosity, say). Sets may be cognitively constituted as object-individuated ensembles in the same way that is attributed to numbers, I suggest (but much further explication and research is needed).

Leaving aside further details and such explorations, the general point is that, *ultimately*, these matters must be taken into account, as part of a complete cognitive account of set theory. They might be relevant to the understanding of *pure* set theory, considering the reliance on the pre-theoretic notion and the familiar real-world examples in the first place. The mathematicians' abstractions and idealizations, the cornerstones of mathematical reality's cognitive constitution, are not *intuitively secure* (as I make a case for here). Understanding how these arise from *concrete* process cases (e.g., the formation of a mental set of the people in the room), a clear account of a well-defined *abstraction-process*, is a necessity for judging their validity. More generally, understanding the cognitive handling of real-world sets, largely automated and even innate, may provide clues regarding the processing of the mathematical realm too, that for the mathematicians brings it about.

8.2.2 Sets as Cognitive Object-Types

Lie up high at the platonic spheres as sets may, for the students (as for mathematicians and for all humans), the actual sets they attend to are always given through a representation. It is from *that* representation that mathematicians abstract and generalize, including to sets that could have no individuated representation (e.g., most subsets of $\mathbb{N}$). There is no "direct access" to the realm of sets. Understanding sets and their scope, accounting for this ontology, could accordingly be achieved only through understanding these representations and their abstractions and idealizations. But whereas traditional "cognitive", intuitionist and constructivist approaches altered ontology in order to fit a simple fixed basis, I aim at the complementary challenge of developing the cognitive-computational basis in order to account for the accepted, standard mathematical ontology. (Ultimately, the "true picture" of set theory would probably be a combination of both, enriching the computational grounding while restricting ontology where it has been taken too far).

Chapters 6 and 7 pursued this dissertation's general representation-first agenda by developing a framework for *cognitive object-constitution* ("objectification") and analyzing in these terms ordinals, cardinals, and the natural numbers. Metaphysically, the approach has left room for mathematics still to be independent of the mind. But it took mathematics and its regularities to apply to cognition itself, and through this, be picked up on by cognition along similar lines to worldly objects and regularities. In this section, I now consider the application and extension of this approach to the object-types that the student comes to master throughout the course.

The fundamental drive of the development of object-types (Section 6.5) is bottom up, forming richer, higher-level arrays (through extensions and amalgamations). Regarding set theory too, the shift of focus from the ontology to its cognitive constitution is accompanied by a foundational shift from the conscious outset to a bottom-up computational view. The conscious level's brand of simplicity (that supported the principle of comprehension) aside, and notwithstanding set theory's great ability to represent just about anything else in mathematics (and in computer science in particular), here we care about the inner workings, which may be hidden by automation etc. Accordingly, this cognitive approach does not take sets for the atomic building blocks of the abstract universe, but the cognitive procedures that bring their object-type about, bottom up. Let us work our way up from the simpler of set representations, then, and develop the topic along the way.

Sets are successfully mastered as an object-type, first and foremost through concrete sets. The generalization of the various representations into a single unifying object-type involves the integration of *translations between them* (e.g., within amalgamations), which may then support the seeming abstraction from the particulars of the representations. Importantly (following Section 6.7.2 and Chapter 7), at the end of the course, the student might *still* have no clue as to the full structure, the domain of all sets (if there is one). The challenge of grasping that full structure *is* the story of the case study.

To set the expectations: Chapter 7 pursued a very particular point, which it could thus analyze rather elaborately. The mathematics covered in the scope of the current section includes that of the former topic, and far more. Accordingly, this section does not aim for a substantial account of the cognition behind set theory, only a preliminary and incomplete sketch.[100] Beyond its value as the beginning of such an endeavor, it could hopefully suffice for providing some basic insight too, which I draw in the last subsection here (8.2.3). Mainly, it is here to serve as background for my account of the illusory intuition, to follow.

[100] In particular, the discussion does not engage with the vast literature on intuitionism and the various constructive set theories with their technical riches.

8.2.2.1 Raw Working-Memory Sets

The items of cognitive representation, the high-level conscious ones that constitute the person's grasp of, for example, mathematical ontology, are those three to four individuated items managed in working memory (Section 5.1). In numerical cognition, this is the core of the perception of the discrete; it is the basic source of numerical judgements: subitizing. From that core, numerical competence is extended through further means (approximation and the procedure of counting, allegedly). Numerosity already assumes some basic delimitation of the domain of objects that are to be counted, though this may not yet be the permutation-independent notion of a set. But just like numerical capabilities (are thought to) extend the elementary case of working-memory items, set theory too, I propose, may extend this extremely restricted case, which already exhibits some fundamental issues. The set-representation at stake, then, is the following object-type:

$\underline{\text{Set}_{\text{WM}}}$: A set that is given through (i.e., represented by) some explicit individuation of few-enough items in working memory, e.g., $\{\text{John}, \text{Barry}, \text{Natasha}\}$

To stress first how degenerate this case is, note that Set_{WM} does not support even a fundamental operation such as general set union (e.g., $\{1,2,3\} \cup \{3,4,5\}$ is not a Set_{WM}).

The fundamental novel operation is the forming of the collection Set_{WM} as a cognitive item, which can now take a single place in working memory. With regard to this item, its (i.e. its object's) elements are not individuated; its slot in working memory holds but one, the set itself. But all of them, in this degenerate type, still occupy the other slots in working memory, individuated. Thus, for example, the mental consideration that John is an element of that Set_{WM} builds on having both him and the set as items of working memory (where Barry, for instance may now be out of attentive focus).

The very depth of set theory's hierarchy over shallow mereology (whose cognitive analysis I omit) might not at all come naturally. The student must come to accept sets such as $\{\}$ and $\{\{\}\}$ and learn to treat the latter as distinct from the former. She must progress from attending to sets of things to sets *as* things, as an operation that can be *iterated*. In effect, the student must assimilate the relevant computational processes. This is like programing a code library that supports the standard type of set and its operations (offering functions such as `Unite(A,B)`, queries such as $a \in B$, etc.). It requires a more complex framework than one that would suffice for non-hierarchical applications (able to treat the very same object both as a collection and as an element of other collections). In this sense, the cognitive object-type itself must incorporate this *computational* richness, in order to reflect the object-coherence of these allegedly platonic sets. The procedure that takes an element x and produces (as an item in working memory) the singleton {x} should be coordinated with the one that produces the element of a singleton to cancel out.

This exemplifies an important and more general point: The computational richness that determines mathematical (perhaps platonic) objects is fully reflected in the representations themselves. At this basic level of set theory, the student with her mind in a sense *implements* the mathematical notion of a set, which can be given a purely computational underlying meaning. The introduction of any object that transcends that (an issue at which we soon arrive) – is what needs to be justified.

Other fundamental properties of sets must also be "programmed in", reflected by the object-type Set_{WM}:

Working memory is a hypothesis with empirical support, not a detailed computational model that is neurologically accounted-for. But assume that the slots are all different, individuated somehow; that there is some consequential difference between holding the items as (John, Barry, Natasha) and as (Natasha, Barry, John). Indifference of the Set_{WM} to the order of its collected items would better be stated explicitly in class (though this may probably be gathered from the linguistic usage of "sets", likely much earlier in life). Indifference to permutation (as already discussed in Section 7.5) must be assimilated into the object-type's object constancy (as part of its object-coherence). Ultimately, these particulars of the representation through which the set is given may become automated and to some extent "disappear", turning the focus to the set itself (a phenomenon which I elaborate in Section 8.3.1).

The more general fundamental facet of a set's object constancy reflected here is that of extensionality; the independence of the set of *anything* but its elements. The independence from the mathematical representation is rather trivial for Set_{WM}, concerning only symmetry (assuming that working memory does not allow for replicas). But the cognitive abstraction includes more than that. Stumbling upon {John, Barry, Natasha} in a different context, at a different time, for example, should lead to its recognition as the very same set. This may be too easy to grasp implicitly (and too standard in abstract mathematics) to warrant an explicit mention in class. But it still cannot be assumed as a default for the cognitive object-type; this is a regularity (Section 6.4) that needs to be integrated into it. Likewise for the symmetry requirement, which also may not necessitate an explicit mention (before the next course comes to formalize set theory upon logic).

Although there is some complexity to Set_{WM}, it might not offer much value in itself. Let us now turn to fundamental enrichments of sets as object-types.

8.2.2.2 Surveyable Sets

Working memory is extremely limited in the number of items it can hold (though not in their richness). The first, most mathematically substantial "idealization" of Set_{WM} is not only cognitively feasible, but also cognitively substantial. Working memory is only the locus of attention; the stream of items that parade through it is central to cognition in general – and to that behind (implicit) sets and (explicit) set theory too. Humans may not be able to perceive each of many objects individually, all at the same time. But their mental toolbox is rich enough for them not to be utterly constrained by this limitation in its particular form. They transcend working, short-term memory, relying on the following:

i. Long-term memory ("Let's see, who are the people I've met in the course of the day?")
ii. Cognitive operation in time (e.g., counting)
iii. External reality (e.g., the toys on the table), including
iv. External memory (e.g., the blackboard)

Humans can usually survey a collection, controlling their attention to pass sequentially through each of the objects individually, ideally once and only once (unless there are too many objects on the table in disarray, for example). The control may be done according to some explicit rule ("scan the table from left to right"). *It is this rule that is an item in working memory*. From our cognitive point of view, it is this rule that not only determines the set, but indeed *constitutes* its representation.

The proposed object-type, then, can be described as follows:

$\underline{Set_{survey}}$: A set that is given through some explicit rule for the sequential individuation of (a finite number of) items through working memory. (The sequential individuation need not actually take place; it just needs to be a potential perceptual action of the object-type).

Set_{survey} relies on and augments Set_{WM}, for which the rule was void. (This sort of relation between object-types awaits further development of the objectification framework). Some fundamental operations build on Set_{WM} in an essential way. For example: The mental validation that some object is an element of some Set_{survey} (as a "perceptual" procedure) requires having both that element and the set as items of working memory (where other elements of the set may be out of attentive focus). The discussed issues and procedures (and their computational support) remain. But now, for example:

1. Two sets can be united. To give an example that may reflect a real-life cognitive operation: A child could notice that there is a Set_{WM} T of (a few) individuated toys on the table; turn her head and attention and notice another Set_{WM} R of toys on the rug. She may then come to realize $T \cup R$: that there are a few toys on the table and on the rug – though no longer all individuated. This, as implicit mathematics, that the brain possibly implements but does not perceive as such (i.e., consciously represent the set as an object). At the beginning of the course, the student is now to reflect upon this and master the object-type.
2. More generally, a family Set_{survey} of sets (also given through a rule, an item in working memory) can be united.
3. For a Set_{survey}, any concrete subset (attended to in life or during the course) can be produced by composing the constitutive rule with a rule for filtering (analogously to the axiom of separation). For examples, "scan the table for toys but include only non-dangerous ones"; or "the numbers below 20 that divide it".
4. The powerset construction can be systemized into a sequential rule.

Through the richness of rules (including the *composition* of rules), finite set theory as used throughout the course is now supported. Set_{survey} tracks (faithfully represents) many familiar notions and their usage.

In general, with greater power come greater challenges (as is a primary theme for us here). For $\mathrm{Set}_{\mathrm{WM}}$, the (reasonable) assumption on working memory was that different slots necessarily contain different items; in other words, that working memory simply avoids duplicates to begin with. However, for the sake of the general expressibility of rules (not just ones that would necessarily select each item only once), the duplicates have to be managed. That $\{1,2\} \cup \{2,3\}$ is $\{1,2,3\}$ rather than $\{1,2,2,3\}$, is not a computational default. This issue may be integral to the general cognitive usage of sets. When perceiving a set of objects on the table by perceiving individuatedly (as $\mathrm{Set}_{\mathrm{WM}}$) those on the left side and those on the right, an object in the middle might be perceived twice. This duplication is a product of the framing of the survey through the rule, and fixing it is necessary for the object-constancy (with regard to other possible rules) of the surveyed set. In programing terms, a duplicates check would ordinarily have to be coded in, which complicates the procedures. Accordingly, in class, although particular sets are usually given through representations with no multiple references to identical elements, the general fact would better be stated explicitly.

Another challenge is for indifference to permutation. Being provided through a rule does not mean that the represented $\mathrm{Set}_{\mathrm{survey}}$ may depend on the rule in an essential way. Indifference to permutation must be extended to permutation of *rules*. For example, the perceiver should realize that turning the table before surveying its objects from left to right would produce the same set.

This point is a most general one. The general *form* of the representation may be essential to the object-type (which for $\mathrm{Set}_{\mathrm{WM}}$, for example, inhibits it from supporting the procedure of set union). But many important implementational "details" of the representation must be managed so as to not influence the represented object itself (as a cognitive analog of the worldly object). Object-constancy means that a represented set object (of types such as $\mathrm{Set}_{\mathrm{WM}}$ and $\mathrm{Set}_{\mathrm{survey}}$) should *ideally* not depend on the particular representation through which it is given (e.g. $\{\mathrm{John}, \mathrm{Barry}\}$ vs. $\{\mathrm{Barry}, \mathrm{John}\}$); just, again, on its form. This is *abstraction from representation* (which, too, I do not develop further, as such an account would require).

Further abstracted, the rule through which a $\mathrm{Set}_{\mathrm{survey}}$ is given may be *non-deterministic*, translatable into a multitude of deterministic rules such that the choice between them is inconsequential to the set they would bring about. Scanning a table for its objects from left to right, independent of the perceiver's angle towards it, would be such an example. This could perhaps be supported by some extended, higher-level $\mathrm{Set}_{\mathrm{Nsurvey}}$ object-type, which can be treated as a $\mathrm{Set}_{\mathrm{survey}}$ in various ways (another awaiting development of the objectification framework).

The ideal of "seeing past" the inconsequential details of the representation is something that the objectification system strives for. Simple and repetitive translations between representations – become automated. In principle (as discussed for $\mathrm{Set}_{\mathrm{WM}}$ above), this should be true[101] for translations between mathematical representations ($\{\mathrm{John}, \mathrm{Barry}\}$ vs. $\{\mathrm{Barry}, \mathrm{John}\}$) as it is for natural frequent activities (assuming GCP, p. 118). However, unlike for $\mathrm{Set}_{\mathrm{WM}}$, much of the related mathematical activity is not at all simple and repetitive. Although the sets $\{30 \leq \mathrm{n} \leq 40 \mid \mathrm{n\ is\ prime}\}$ and $\{31,37,39\}$ are actually one and the same, this cannot simply be read off of these representations. Perceiving them as the same passes through rare-enough particular calculations that require conscious control. At the far end, such comparisons can actually stand for (e.g. code) open problems. Learning, assimilating an object-type, need not necessarily imply an automatic, attention-free mastery. The assimilation may be "stuck" at the pre-automated stage – and necessarily so for cases such as the latter. Assimilating what is required for the object-type and its coherence, namely the computations that must be performed in order to handle the object, works if this is kept at an explicit conscious level of knowledge too. Further acquaintance with the subject matter, through following proofs that concern the object-type and mostly through solving exercises, can then automate parts to produce a mastery. (The value of exercises, to speculate, may be in the way the objectification system comes to master manipulation procedures *qua manipulations*).

Just to make sure: $\mathrm{Set}_{\mathrm{survey}}$ is rich enough to raise the issue, but these remarks on representations are completely general.

[101] If mathematical cognition follows the general ways of cognition for this too (the central working-hypothesis, p. 131).

8.2.2.3 Linguistic Representations

A general type of representation, which is particularly important for this case study, is that of a *linguistic* representation, as in "The students in the classroom". I suggest that it originates from extensions of object-types such as above. Set_{survey}, say, is extended into $Set_{surveyL}$, which offers extra procedures that are linguistic in nature. As always, these have to be properly coordinated with the Set_{survey}'s original procedures. Linguistic set union A∪B, for example, is coordinated with the formation of a "unified", concatenated procedure that can survey all of the constituents of A and then B. The result comes out the same set of $Set_{surveyL}$. This allows the student to pursue valid logical constructions without keeping to the underlying constructive grounding, knowing (non-constructively) that it indeed has a constructive interpretation to switch back to. The linguistic representation of $Set_{surveyL}$ is taken to come, inherently, with a procedure for translating it into an enumeration. It is still grounded in the *potential* for an individuating survey through working memory. In line with non-mathematical language, logical quantification is not constitutive, but rather grounded in a domain that is constituted independently, such as through some survey procedure.

With linguistically imbued set object-types such as $Set_{surveyL}$, it is their *abstraction*, their *detachment* from any surveying procedure, that then produces the reduct object-type for what we consider a class; a *purely* linguistic representation.

8.2.2.4 Denumerable Sets

So far, the restricted represented parts of set theory could be implemented in full, computationally. The student could in principle manage her own implemented object-type as an informational "copy" of the discussed, allegedly platonic set. *Infinite* set theory, on the other hand, is paradigmatic of that which transcends computation. Accounting for it in computationally grounded representational terms becomes much more challenging and controversial. Here, the true tension with the cognitive foundational perspective lies. (But, again, Benacerraf's problem of the mathematicians' perception of mathematical reality stands regardless of one's philosophical position).

The prerequisites are, first of all, a mastery of finite ordinals, cardinals, and their amalgamation into natural numbers as an object-type (Chapter 7). This includes, in particular, the successor procedure: the object-types' general support for producing (as an item in working memory) the successor of a given ordinal or number (as an item in working memory). Upon this procedure, somehow, the student furthermore should grasp *the* natural numbers, as a completed domain (which does not by default come with a mastery of the object-type itself; Section 6.7.2). The domain is not yet a *set*-domain. Though holistically grasped, it may be inherently ordered and unsupportive of completed, infinite permutations. Actual infinity may still be only implicit, in the background (constituting potential infinity in the foreground). It must furthermore be grasped, explicitly, as a meta-procedure, for running some procedures (such as the successor procedure) an actual infinity of times. Importantly, this might not necessarily require the infinite set itself; ω as an order-type (which, recall, might not be based on an underlying set) suffices.

With these in place, the extra step then is in forming the natural numbers into $\mathbb{N}$, as a *completed set*. As such, $\mathbb{N}$ could never literally be embedded as a whole within any mind or computer. But this was already an issue for $\text{Set}_{\text{survey}}$, as it transcends the far stricter limitations of working memory. The basis is the same: The set is given through a rule for the potential individuation in working memory of its elements. This is the fundamental form of its representation. A central difference here is that this potential must *necessarily* stay a mere potential ("necessity", as opposed to impractically large finite sets). This difference can affect the object-type itself (for example, with respect to the procedure that produces a 1-1 mapping with a proper subset, something the student quickly learn of).

$\mathbb{N}$ is the archetype. But the added machinery in place interacts with the previous support for the composition of rules etc. This can produce many concrete subsets of $\mathbb{N}$ (e.g., the prime numbers), as attended to throughout the course. This more general object-type, Set_{E}, can now support interesting bits of naïve set theory (e.g., $\mathbb{N} \cong \mathbb{Q}$).

This barely scratches the surface of the topic and the issues it raises. But hopefully, it is a reasonable sketch of what is going on inside the student's head. To take the wider philosophical perspective:

With cognition sufficiently idealized, any particular number may conceivably be represented. Even then, though, they cannot *all* be represented simultaneously yet in an individuated fashion. The representation of $\mathbb{N}$ itself, as a completed whole, cannot "literally", isomorphically, reflect the elementhood relation with regard to the representation of individual natural numbers. The representation is inherently *meta*mathematical, constructed upon objects that were not explicitly referenced as part of the mathematical language as formally conceived. The mathematician's toolbox (and cognitive representation) also includes the successor procedure and other (perfectly mathematical) procedures and rules as above, which are not sets (even if essential aspects can be coded as such). A cognitive foundation reveals a different order of dependence (e.g. creation) between entities than the set-theoretic one. As cognitively represented mathematical objects, 16,947,233 does not precede, does not participate in constituting, $\mathbb{N}$.

8.2.2.5 $P(\mathbb{N})$

With $\mathbb{N}$ in hand, $P(\mathbb{N})$[102] would be a basic product of the following:

1. The general notion of a subset.
2. The ability to harness the former into collecting all such sets.

But their meaning is not constituted in sheer logical language. They should be accepted as general only if the relevant object-types do in fact support them. This is indeed the case for their origins are in the finite (i.e. $\mathrm{Set}_{\mathrm{survey}}$). Any subset of some finite set F can be produced or recognized, just like we can bite an apple or recognize that it has been bitten. Relations among objects of an object-type are standard material for the objectification system (if mostly outside the limited scope of Chapter 6). Reflecting on the object-type itself in order to collect all subsets of F is not as straightforward or naturally supported (as a limited case of complete domain grasp, Section 6.7.2). And indeed, a method for that is not initially evident for the student. Still, it can be done; a rule for surveying F that can be completed can be turned canonically into a rule for surveying all of its subsets $P(F)$.

With $\mathbb{N}$, the situation is very different, of course. Set_E can represent many concrete subsets of $\mathbb{N}$. This domain (if there could be a well-defined one) would be $CP_E(\mathbb{N}) \subset P(\mathbb{N})$, the "Cognitive Powerset" – any particular subset of type Set_E that the

[102] The general powerset construction is not as obvious as it is usually taken to be. A fascinating historical fact is that in Cantor's own last word on the issue (in a letter to Hilbert), he "state[s] forcefully: the argument that the power set… is available because S is available is illusory" (Ferreirós, 2007 pp. 448-449).

students may encounter and come to represent. However, mastering an object-type does not entail the perception of its domain, so even $CP_E(\mathbb{N})$ need not be consciously available. And most subsets (of $\mathbb{N}$) in fact can have no cognitively accessible rule to represent them. These can never enter mathematical practice; not individually. Yet mathematicians still postulate their existence and take their reasoning to pertain to them too.

Are they generalizing from the subsets that they do stumble upon? Do they reflect, meta-cognitively, upon the Set_E? If so, then that would cry out for cognitive justification that goes beyond introspection: The object-type through which mathematicians approach the subsets depends on the form of its representation, of which they need not be fully aware. This form may have implications for the object-type that are not shared by subsets outside its domain (as almost all of them are). Those latter, postulated subsets are grounded in a variant of actual infinity that allows for an independent action (choice) at every point in the infinite series (which is more than is required for $\mathbb{N}$ itself). The representation itself is thus idealized, beyond its finite cognitive grounding in predetermined rules that produce completed wholes. This is the general notion of a subset (of $\mathbb{N}$) – also grounded in cognitive representation somehow – and the power-set construction then builds on *that*. $P(\mathbb{N})$ collects the subsets of $\mathbb{N}$ that *could* be individuated using actual infinity – in a manner that extends, coordinates with, the subset of those subsets that do have finite representation ($CP_E(\mathbb{N})$).

This discussion only begins to raise the questions that need to be addressed. What I would like to stress here is that, even in idealization, the conception of these sets is grounded in processes. Crucially: the collection of the subsets of $\mathbb{N}$ as given through idealized, infinite representations, comes out the same, *extensionally*, as the collection of subsets of $\mathbb{N}$ idealized *from* their representation, from how they are given. (Through both, any number can be included or not). However, intensionally, and cognitively, these are not the same – and their further idealization need not come out the same extensionally too.

8.2.2.6 Constructions of Ordinals

Fig. 17 Illustration of the ordinals up to ω^ω.

Original: *(Wikip-Ord, 2018)*. Tattoo: Shai (Deshe)

Abstractly defined, abstractly *handled*, the class of ordinals, as a whole, serves as the infinitary backbone of the entirety of higher set theory. Many efforts in the basic course may be put into defining them the von Neumann way and proving the lemmas required for establishing their properties and applications. However, much can be done to endow the student with some intuition, the way intuition is usually developed: through concrete models, even visually illustrated. That this can be done, that a cognitive grasp on a substantial initial segment of that backbone can be given, calls, too, for a representational account grounded in computation. Under the assumption that mathematicians can rely on a finite representation of $\mathbb{N}$ itself, many countable ordinals can then have *essentially finite* representations (p. 116) too. They are thus in a sense finite, constructible, concrete entities. Moreover, these entities need not even be conceived as sets (Section 7.4); they can stand for an order-type or the length of an (idealized) iteration. Let us try to approach them as such, then. [103]

Our initial finitistic computational ontology, which is sufficient for representing the finite ordinals, consists of a single ordinal called '0', an ordering relation '<', and a single action called '+1' which, applied to an ordinal, produces a new, distinct, larger next ordinal.

An entirety of a set of ordinals is *not* specified and is not considered as the ontology here. The infinity of finite ordinals can be kept external to this current discussion, which

[103] See also (Sloman, 2002), on the "visualization" of such infinite structures.

is strictly *finitistic*. To support this, a fundamental move is to consider ω, the first ordinal larger than *all* finite ordinals, as the first ordinal larger than *any* finite ordinal instead.

Classically conceived, we could think of ω as represented by the 2-tuple $<1,0>$, which is lexicographically greater than any finite ordinal n, as represented by $<0,n>$. However, this would require us to extend the underlying ontology later, to 2-tuples, 3-tuples, $\omega+1$-tuples, etc. Instead, let us yield on ordinals themselves (except for 0) as preexisting and maintain the following:

1. An ontology that also allows for *actions*, which always produce new, distinct, latter, next objects.
2. A single principle: for every object of this ontology (ordinal or action), there is an action that produces the successor object.

Let us see how this would work.

For the production of the finite ordinals, the action promised by the principle is the '$+1$' we have already met.

Applying the principle to this action produces a new action, '$+\omega$'. Being the successor of '$+1$' means exactly what it should, namely that the ordinal $\alpha+\omega$ is greater than any ordinal that can result from applications of '$+1$' to α and is the smallest such ordinal. ω is thus not defined relative to all the finite ordinals, but rather to just '0' and '$+1$'. Applying '$+\omega$' to any $\omega+n$ results in $\omega\times2$, and then $\omega\times3$, etc. So, we can reach any ordinal smaller than ω^2 in a finite number of steps (and thus illustrate it pictorially).

Now, we can apply the principle in turn to '$+\omega$' itself and get the successor action, '$+\omega^2$' (the "limit" of $+\omega+\omega+\omega$ ...). For example, ω^2 is defined, as before, relative not to all the ordinals smaller than it, but rather just to 0, '$+1$', and '$+\omega$'. All of these actions, when conceived as classical functions, as a type of correspondence that relates to the whole infinite ontology, are infinite. However, they need not necessarily be conceived as such.

Applying the principle over and over again produces '$+\omega^3$', '$+\omega^4$', etc., and thus all of the classical ontology up to ω^ω. The branching[104] complexity of this whole structure, although far transcending that of the linearly ordered natural numbers, is inherent to these ordinals, not just an artifact of this representation of them (even if in Fig. 17 it was "linearized").

With a machinery of this sort in place for the students, we can revisit ordinals *qua sets*. $\mathbb{N}$ (as a set) and $P(\mathbb{N})$ were taken to be constructed using the underlying conception of actual infinity as a process. Likewise, understanding ordinal-length iteration as a process can *lead* to the construction of sets that encode it (à la von Neumann). But the classical consideration of Benacerraf (Section 1.4) extends easily to the infinite ordinals: this encoding is not a faithful representation. The cognitive representation is different (and I suggest, procedure grounded, dynamically). Metaphysics would do better to follow in its footsteps rather than lose anything to reinterpretations that take us further away from the essence.

8.2.2.7 The Axiom of Choice

As we have seen with $P(\mathbb{N})$, the acceptance of $\mathbb{N}$ (and its expression through the axiom of infinity) is not enough to decide all fundamental matters of infinity (even in that first course, before getting into large cardinals). Left still is the pinnacle of set theory's alleged non-computability (non-constructivity) – the axiom of choice (AC). Formally, it is outside the scope of this case study, as comprehension does not entail it. But it deepens the metaphysics of this cognitive approach and illuminates the difference between ontology, the mathematicians' access to it, intuition, and language.

AC states (in Russell's form) that a Cartesian product of non-empty sets is non-empty. To bring in his own picturesque example of infinite collections, of pairs of shoes and pairs of socks: "From each pair of shoes, choose the left one" is a simple rule for how to choose one of each – across the board. Unlike pairs of shoes, pairs of socks come with no extra properties to differentiate between any two by – nothing but the definitional non-identity of "there are two". The latter alone does suffice (by the laws

[104] There is a denumerable number of actions, where from each one, the potential applications of all lower actions branch out.

of logic) for choosing a single sock from any particular, individuated pair of socks. What one cannot do, lacking a differentiating property, is provide a rule for a *global* choice for all of the pairs of socks at once.

AC concerns non-constructive *entities*, making room for them within the ontology. There are certain supposed types of objects that mathematicians would like to get a handle on (as particular elements, that are individuated), even when they cannot construct their representation and thus witness their existence through logic and continue from there. The underlying logic can thus influence the issue. But unlike for constructive mathematics, which rejects non-constructive *proofs*, the revision of classical logic is not required.

The question, then, is as follows: *Why* might it seem to mathematicians that there needs to be something there, whether or not they can access it through logical language?[105] The historical controversy around this is well known; probably no answer could be accepted across the board. But I offer one that is not only reasonable but integrates naturally within our greater cognitive story.

The set-theoretic ontology itself is static. Some sets represent functions (in the modern sense of the word), but they are not operations in any temporal sense. Function-*applications* are metaphysically outside the scope here. They are concealed by their non-faithful representation. Application as a primitive notion can relate to the logic: Given any set-expression s, one can add "{" and "}" to get "{s}" and by iterating this for ∅ "produce" more and more Zermelo numbers. These are operations on linguistic representations, which can be idealized to continue potentially to infinity. Yet standard, first-order logic enforces a particular flavor of finitism (with its well-known limitations or shortcomings, as discussed in Section 1.1), in which the expressions themselves are inherently finite, putting a limit on what can be represented. (This is why Zermelo's construction could not be continued to infinite ordinals).

[105] "This was particularly objectionable to mathematicians of a "constructive" bent such as the so-called French Empiricists Baire, Borel and Lebesgue, for whom a mathematical object could be asserted to exist only if it can be defined in such a way as to characterize it uniquely". (Bell, 2015)

For the typical "pure" mathematician, who is not logically and foundationally inclined, not attributing any significance to logic's particular form of finitism and to the linguistic representations of the reality with which she deals, all this, at least nowadays, is usually a non-issue: There is no apparent problem in choosing (as an act) over referring (which is also an act) here. As long as infinity is granted: if she can be given (a finite description of) an infinite set of non-empty sets, then the very same idealizedly infinite reasoning that leads her to identify each one's non-emptiness, all at once, allows her to identify an element of each one, all at once. The fact that, as a logical framework would have it, she is allowed to move from "identifying non-emptiness" to "identifying an element" finitely many times but not infinitely many is *an artifact* of logic. This is a distinction for which that typical mathematician does not care, and in fact even needs to be trained for explicitly in order to assimilate. The issue can be quite sweepingly missed by the students at the end of the course, *even when they have mastered proofs using* AC (equivalence to other principles, or applications to other mathematical areas, as in "Every vector space has a basis"). This is not to say that such a distinction can carry no meaning, just that it is a finer, ancillary one – and one that is not constitutive of the ontology as the mathematician perceives it. That those undefinable entities, which transcend all that a logical language can handle and are more complex than all that the mathematician actually meets, grasps concretely, may thereby also carry surprising consequences – is *unsurprising*. Her intuitions are formed about what she *can* and does perceive!

In this regard, it is illuminative to consider a related, also quite common misconception (even among professional mathematicians) – that AC is not needed for *countable* collections. That "pure" mathematician, not focused on the finiteness of logic (especially as set theory is so much about the infinite) seems to have an underlying intuition here that dictates the following: discrete processes are something she understands, perhaps even "perceives" (as has already been accepted with $\mathbb{N}$). Idealizing an action as being taken at every minute ad infinitum might not seem to go far beyond idealizing the infinitude of minutes themselves. Where actions *do* seem to lend themselves much less to idealization than mere discrete times is in complete immediateness, not isolated in any temporal segment of their own (as working memory surely requires for individuation).[106] Making choices at each and every point of a continuum, no separation between them, may seem more like an abstract definition that one cannot truly perceive than an idealization that is in the same region as those of very quick choices or an endless sequence of them. Thus, AC_ω (choice for only countable collections) is rooted in processes that we can in a sense conceive of or "perceive", whereas AC in general is too abstract, devoid of procedural content, and draws its appeal from a simplistic, unifying linguistic framework instead. This can receive further support by considering that AC_ω is an intuitively interesting midpoint of ontological commitment: Overarching theoretical elegance aside, many of the actual results for which mathematicians would like to have AC only require AC_ω, whereas many of the counterintuitive results of AC (which raise suspicion against it) cannot be derived from AC_ω alone. This, too, conveys how ontology as represented in cognition is founded foremost upon processes, rather than on their linguistic products and the laws they uphold. It, too, demonstrates the danger in the appeal of succinct, elegant linguistic theorizing. AC may conceal different procedural groundings in different cases it is presumed to apply to, just like "the power-set of" may conceal a very different nature when applied to finite sets, $\mathbb{N}$, $P(\mathbb{N})$, or unspecified, general abstract sets.

Let us conclude with a cognitive (yet realist) metaphysical reading along the lines of MAC (p. 59). Consider again Russell's picturesque example through shoes and socks (just two denumerable sequences of pairs, assume for concreteness). The situation can be viewed as a two-sided game between the presenter of the challenge and the solver (who is to find the required global choice function). The presenter used a finite number of words to present the situation. He did not specify in what way the pairs are different. *Conceptually*, then, he necessarily “replicated the same pair over and over again” (relying on a single represented object). Indexing the pairs by the natural numbers was just a formal means for giving grounds to the content-less claim that they are actually different. As for the solver, she is allowed to choose a single sock from the first pair

[106] But: "Zermelo asserts that ‘the purely objective character’ of this principle ‘is immediately evident.’ In making this assertion Zermelo meant to emphasize the fact that in this form the principle makes no appeal to the possibility of making ‘choices’". (Bell, 2015)

(simply because it is a pair; this is part of what the set object-type supports). However, assuming that the game is "fair", that both parties are acting within the same systematic framework such that the presenter is not given an advantage in abilities, then the solver could analogously repeat, replicate her choice over and over again, by simply indexing the same choice for each identical but differently indexed pair.

For the platonist, the mathematician's role in mathematics is in perceiving and proving, reasoning about it. It is not a game between two equal-leveled players but a one-way connection of mathematicians as solvers trying to refer to and reason about the boundless, representation-independent mathematical reality. There simply *are* (upon postulation) infinitely many pairs of socks, *truly* distinct rather than a replication of a single one, even if they cannot perceive them as such. Mathematicians can *refer* to a whole infinity of them, even if this infinity is *in no sense* finitely specified. And so, an essentially different choice has to be made for each pair.

This is the metaphysical leap from which MAC) would probably refrain. According to it, mathematicians do *not* reason directly about reality "itself", but about cognitive representations of reality. This cognitive representation is the former "presenter" player, whose resources may ultimately thus be just as finite. Simply saying "an infinity of different pairs", different in no sense other than by sheer postulation, is a linguistic construct that does not actually refer to anything. It is merely schematic; it can be applied to many entities; to various cognitive representations that fit its mold. Anything it applies to (say, identical pairs differentiated only by numeral indexes) must be an actual, *essentially* finite representation. And for such an actual entity, in a fair game, AC is but a tautological truth[107]. Alternatively, there could be valid reasons for a division into sides that includes a disparity between their idealized capabilities. But this could still be a cognitive in-house matter (e.g., some architectural consideration), which would best be explored in such terms.

[107] To bring in the words of Bishop: "A choice function exists in constructive mathematics, because a choice is *implied by the very meaning of existence*". (1985)

8.2.3 Interim Conclusion: General Skepticism

What is the value of the intuition for the existence of U? Modern cognitive science generally forbids naïve reliance on the way things consciously appear to be. Let us here consider from that point of view the computationally grounded representation of set theory (along the lines of the sketch above), to elucidate the skepticism more concretely.[108] I personally do not regard this traditional skeptic role of philosophy to be particularly interesting for this case. But it does have to be spelled out, as the philosophy of mathematics is very much behind in accommodating for this modern scientific view. And mostly, it rests on our path towards a more constructive (if suggestive) account of what might actually be going on.

The procedures on offer by the object-types (which I take to stand at the heart of the representation of the ontology) are consciously available to the student as things that can be done or perceived of sets. However, many of their inner workings, the internal computations and assimilated regularities that go into managing and indeed bringing about these object-types, seldom received the students' direct attention, not to mention being expressed linguistically, and even more rarely formally. (We have seen a detailed example concerning finite ordinals and cardinals in Chapter 7). As the procedures are run over and over again, the details may "disappear" into automaticity, if they were not automatic to begin with, brought about innately (as in Section 8.2.1). The more essential, the more frequent they are – the more invisible they become (or already are). And this includes a vast apparatus of operations already in place, an apparatus that has been serving the students in their daily lives since childhood. (As with Piaget, who famously explored how logico-mathematical abilities appear at specific ages in early childhood – a long way from the grown-ups now studying set theory).

Generally speaking, the conscious level is the highest cognitive level, referring to the most overarching complexes of operations. It does not provide a mathematically precise (if generalized) expression of the complex inner-operation. It is at best a heuristic

[108] Lacking a full account of naïve set theory in terms of the objectification framework (something that would have to further be developed too), the following discussion is necessarily kept to general terms.

approximation. It is rarely fully-detailed, and never with any subjective yet *valid* assurance of being so.

Consider, then, the student at the end of the semester, after having successfully mastered the various partial object-types (including some amalgamations and other interactions between object-types), and perhaps even some completely general object-type of *set*. Say she now contemplates how the prevalent notion of set, or rather a guideline for sets' existence, could be made precise. Part of what she is after (perhaps non-consciously) is *simplism* (intimately related to *elegance*). Roughly, this means involving as little *conscious* (and in particular subject-specific) detail as possible, together with as much generality as possible, i.e. reigning over all that needs to be subsumed under the notion, ideally. Information-theoretically, what she seeks is a type of *compression*, from the concrete sets which were met and represented throughout the course into a singular notion from which they can easily be gotten, and that supports the construction rules and related notions. And, if indeed there was a coherent object-type of *set* mastered, this would be to characterize it – at least extensionally. This particular sort of compression, however, is not to be easily related to those familiar, mathematically explored ones, where everything is explicit, 'conscious', single-layered. Rather, this *mental* compression is into *conscious* simplism, which makes use of i) pre-existing machinery, which renders complicated but neurologically supported operations informationally cheap, and ii) any preexisting data and gained experience (whether purely mathematical in nature or not, which puts the mathematical issue into the broader context, which also includes its applied side, to real-life reasoning).

The paradigmatic *mathematical* road to such compression is through abstraction. But *cognitively*, compression poses some central risks:

- As it turned out (see *categorization*, p. 140), the structuring of our conceptual space in general is not given through definitions as intersections of properties (a complication that arose naturally for my objectification theory too, p. 157). It is not clear, then, that the cognitive compression could be related to a classically, axiomatically defined mathematical conception, at least not in a natural and meaningful way.
- Cognition is not shaped to value *absolute* precision for its own sake (i.e., "lossless compression"), which can be a very costly overkill. Cutting corners heuristically is

the general rule (i.e., so-called "lossy compression"). Furthermore, the distribution underlying heuristics' "working *most* of the time" is purely pragmatic, made to fit those cases that living beings are most likely to encounter, or even just the most likely to *have* encountered in the past (as was in a sense the point of Chapter 67). In our set-theoretic case, care should be taken to focus on the set formulas that are actually encountered in the naïve course. If, instead, the weights are set according to an independent logic-theoretical consideration (namely, treating all well-formed formulas equally), then comprehension might actually collapse (name proper classes) for many such usages. Language makes room for *over*-abstraction.

The preliminary skeptical conclusion from all of this: We need not have waited for an obvious paradox such as Russell's to arise in order suspect the faulty principle of comprehension. Its intuitiveness may spring from other sources than just pure-and-simple mathematical truth, sources of cognitive efficiency with dubious mathematical legitimacy. At this level of generality, though, this skepticism applies to other intuitions as well. Why should we have suspected comprehension *in particular*?

8.3 Omniperception

Mathematics, presumably, has no paradoxes. Mathematical reality is what it is, just like physical reality is what it is. It is only the *perception* of (either) reality that can go astray. This suggests a general approach of considering paradoxes as cognitive illusions in the appropriate semantic realm and exploring them as such. This is what I now pursue in order to transcend the general skepticism about the intuition for the universal set and instead offer an account of the cognitive illusion that stands behind it. For this, I employ the "general cognitive principles" strategy (p. 119):

In Subsection 8.3.1, I introduce the phenomenon of omniperception, which is central to visual perception. I introduce some central examples to this general phenomenon, while formulating its presentation in the most general cognitive terms, as a possibly more general phenomenon (than just vision), to be considered explicitly in Subsection 8.3.2. In Subsection 8.3.3, I show how the phenomenon, if it indeed takes place in the cognitive representation that underlies naïve set theory too, would produce the intuition for the universal set.

The usual GCP caveat applies: The familiar cognitive topic is used for an analogy in order to help guide our speculation as to what may actually be happening. Future scientific work will have to validate the picture and either refine it – or refute it.

8.3.1 Omniperception in Vision

Consider the route the data take in a person's (conscious) visual perception of a scene. The starting point of the data is all that stands (as objects in the world itself) within the field of vision (not hidden) and that is visible in principle. The amount of data is practically unbounded. Moving ahead (and adding some limitations), the *cognitive* starting point of the visual data's journey is the person's retina, which still contains more than 100 million photoreceptors. In contrast, at the far, *conscious* end (which is not the only end to perception but is the focus of my analogy), lies a very high-level, integrative and even holistic perception of the scene, with working memory's extremely limited capacity for holding very few items at a time. As a representation of the scene, there is a large *gap* of informational content between the worldly starting point and the consciously perceived end. The actual details of how the mind works out this gap in essence *are* what the cognitive science of vision is about (though some concerns non-conscious perception). Alongside this matter stands the question which rests at the heart of my account of the universal set: *To what extent* are we aware of this gap?

The one-sentence answer would be: Hardly at all. The cognitive machinery is necessarily an expert at hiding most details of its own workings from the conscious level, which is thus freed to concentrate its extremely limited resources on the issues that cannot or ought not to be handled automatically, but rather require conscious attention. Thus, phenomenologically (in the cognitive science sense of "what it is like", not in the loaded philosophical one), we make a *leap*: it seems to us as though we simply (directly, unmediatedly) perceive what there is in front of us to be perceived (leaving room for refined details or noticing

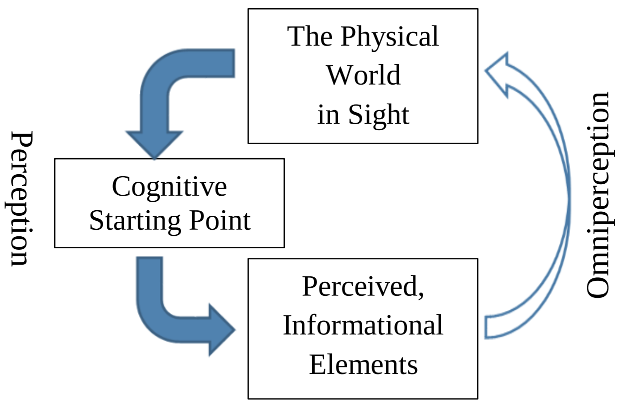

Fig18 . The full arrows represent the flow of information, where the empty arrow represents the experiential reversal.

through attention). The cognitive system transparently mediates our active place in the world.

This tension, between the phenomenology of perception and its actual workings, is a central theme underlying much existing scientific research and one which has brought about a mega-conception that arises from a great many varied empirical findings. Accordingly, it has been titled as a "*Grand Illusion*" in (Noë, 2002)[109]. Noë's own formulation of the phenomenological side in question is the ***Snapshot Conception***:

> According to a conception of visual experience that has been widely held by perceptual theorists, you open your eyes and —presto! — you enjoy a richly detailed picture-like experience of the world, one that represents the world in sharp focus, uniform detail and high resolution from the centre out to the periphery. Let us call this the snapshot conception of experience. (Noë, 2002 p. 2)

This formulation offers the general picture, so to speak. It does, however, miss the third spatial dimension, the temporal dimension of visual perception, and more. Here is another, less committed popular exposition of this tension between the phenomenology of perception and the actual system workings:

> We all feel we are conscious of a rich and detailed world in front of our eyes. Yet outside the dead center of our gaze, vision is amazingly coarse. Just try holding your hand a few inches from your line of sight and counting your fingers. And if someone removed and reinserted an object every time you blinked (which experimenters can simulate by flashing two pictures in rapid sequence), you would be hard pressed to notice the change. Ordinarily, our eyes flit from place to place, alighting on whichever object needs our attention on a need-to-know basis. This fools us into thinking that wall-to-wall detail was there all along--an example of how we overestimate the scope and power of our own consciousness. (Pinker, 2007)

Despite its eminence for the working cognitive scientist, the conception that has been alluded to lacks a theoretical framework on which I would have liked to build the

[109] Which is dedicated to that conception, mostly from a critical, non-representative point of view.

forthcoming analogy[110], one in which to fit and explore the scope and power of our own consciousness (the *gap* from reality) – and our overestimation of it (the *leap*). I thus offer here a rudimentary framework of my own for the general conception that I term *omniperception*[111].

For the information, getting all the way from the richness of the worldly starting point, through the cognitive entry point, and into the confines of conscious perception, is achieved chiefly through (a combination of) the following:

i. Constructing higher-level object-representations over multitudes of those basic data entries.
ii. Throwing away data (cleverly).
iii. Guessing (e.g., based on previously learned regularities, as in Section 6.4).

These are interdependent aspects of the process, not excluding techniques. (i) (which I have attended to in Section 5.3 and Chapter 6) presents no phenomenological discrepancy (of the sort that (ii) and (iii) can and will). Higher-level objects are the very contents of our phenomenology, too. We immediately recognize a friend's face; the fact that it may be under a somewhat different lighting than usual or from a slightly different angle, and so with much variance at the atomic informational level (which we may not even take notice of, i.e. (ii)) – would not surprise us. *High level* means capturing the underlying aspects that matter to us most; that the system compresses away most of what does not, we take for granted (even if are not aware of the computational complexities behind *object constancy*). The informational content that does make it all the way to conscious working memory, although not manifested in many different items, is contained in the richness of the few items. And ideally, if the object-types are sufficiently rich or high-level (i.e., with the richness integrated into the object-types), no information loss is necessary at all. We cannot perceive an array of seven independent co-interacting actors on stage as such. But if we could perceive a few of

[110] With respect to visual illusions, for example, "There is no better indicator of the forlornness of this hope [the hope of some to find a general theory] than a thorough review of the illusions themselves" (Robinson, 1972). (Since then, a wide variety of illusions has been adding up much faster than their generalizing theories).

[111] This term avoids Noë's restrictive "snapshot conception", while capturing the heart of the matter, in preparation for the generalization beyond vision.

them as one (e.g., chunk the romantic couples among them into harmonious entities perceived as one), thereby reducing the overall number of objects to four – then conceivably we could. The limitation to but a few individuated objects is no limitation at all unless the objects themselves are also constrained.

But they are. Objects are objectified as such in order for us to be able to generalize about them, and through this generalization, make predictions and control them. A ball would for the most part act in a similar way at different times, in different locations, against some surfaces, etc., and so perceiving it as the same object or type of object facilitates interaction with it. But this requires that such generalizations about it could be made, and furthermore involves *prior acquaintance* with balls on which these generalizations could be built (into an object-type). The extra informational dimension of what is new to us, not assimilated into the object-type, is what we need to attend to, consciously. This is the reason why we need to attend to a visual scene and perceive some (smaller) objects in it and the relations (interactions or affordances) between them, rather than just perceive the scene as a whole as one mega-object.

Thus, the limitations on our consciousness are too severe, or rather our high-level concepts too rich and refined, for the story to end with (i). First of all, a single scene typically contains many (of what are for us) high-level objects – more than we can handle. This means that we have to actively select what to attend to. Secondly, and crucially, the objects themselves are not simply given to us, to our perception. Much of the cognitive work (including data-reduction and data-completion – throwing away and guessing) goes into determining and constructing the object-representations to begin with, which is computational work that the perceiver is not aware of.

In particular, the *cognitive* entry point (in between the world and its phenomenological perception) is, in principle, phenomenologically irrelevant. To explain this through an extremely common misframing in the literature: It is common to classify “Terror Subterra” (Fig20 .) as an optical *illusion*. The reasoning is as follows: The two monsters take the same number of pixels on the screen, but more importantly, determine an equivalent angle on the cornea. Yet we cannot help but perceive the higher monster "in the back" as much larger than the one below, even though they are in fact perfectly identical. The fallacy is that *there is in fact no such fact of the matter*: not for the indicated objects themselves nor for the suggested monsters out in the three-dimensional world (as opposed to the figures) – which is what the visual system is out to reconstruct. A larger, proportionally more distanced real-world monster would in principle hit the cognitive starting point in just the same way. The visual system *guesses* the worldly position based on extra context – in this case, the locations within the perspective tunnel (which suggests very different distances and sizes). To insist that it is still an illusion is to expect the perceiver to be able to take the extra (conscious) evidence (that it is all on a paper) *top-down into the perception process*, and through it, override the "illusive" background information, in order to attribute a different 3D depth. But such conscious control is not ordinarily required in real life, outside of such artificial tasks. Even when something depicts (e.g. a television), it works for us precisely because we see through it as a medium, cooperate with its attempt to masquerade as a real-world object (that is further away). And so, 3D depth attribution is better automated (be the mechanisms involved innate or acquired early on). Similar considerations apply, for example, to the “Checker shadow illusion” (Fig19 .) (alongside many others of this sort). The general point here is that all of this, regarding figures on paper or on screen, applies to the retina just as much (i.e., for two

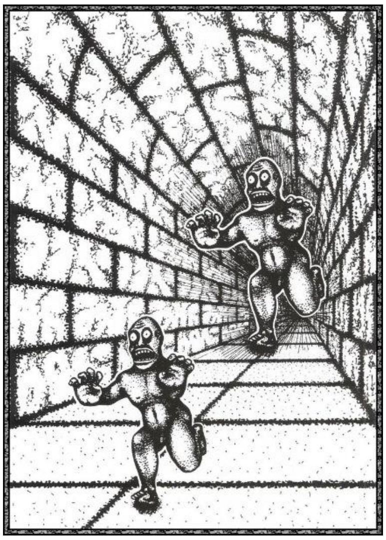

Fig20 . The monsters are identical in size. Terror Subterra / Roger N. Shepard (Mind Sights, 1990).

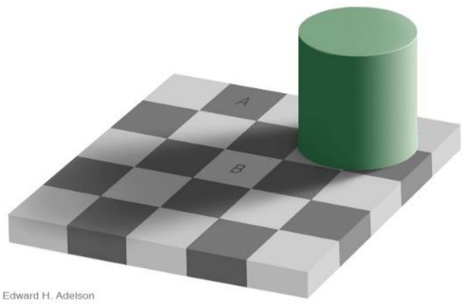

Fig19 . Squares A and B are of the same color. Checker shadow illusion / Edward H. Adelson (1995)

monsters that happen to hit the retina as the same size). To the (non-meta-cognitive) phenomenology, the particular *cognitive* entry point is in principle irrelevant. It mediates and constrains the transfer of information from reality to consciousness, but the information about *it* is non-informative of what there is to be perceived (Fig22 .). The *experience* ideally ought to be one of *directly* perceiving (and being in) reality. Thus, we can simplify the diagram for omniperception (Fig18 .) into Fig22 .

Reality itself is complete. The experience thus transcends many limitations of both the cognitive entry point and the cognitive inner workings. The visual system automatically “fills in” information to reflect a complete model of the actual (visible) world. It need not – and most likely does not – actually reconstruct a complete model. As Dennett has pointed out, "The absence of representation is not the same as the representation of absence" (1991). With respect to the role of objects in perception, we can break down this “completeness” into two aspects.

The first aspect is the “completeness” of each object (in spite of its partial perception). For example, physical objects are represented and experienced inherently as three dimensional, even though the 3D structure is only guessed using a variety of methods from lower-dimensional inputs. As usual, the system does so using not only the partial bits that it has actually acquired of the scene, but also by relying implicitly on learned regularities (including those learned by evolution, i.e. innate). This aspect would be expected to be relevant to the “perception” of mathematical objects such as sets too.

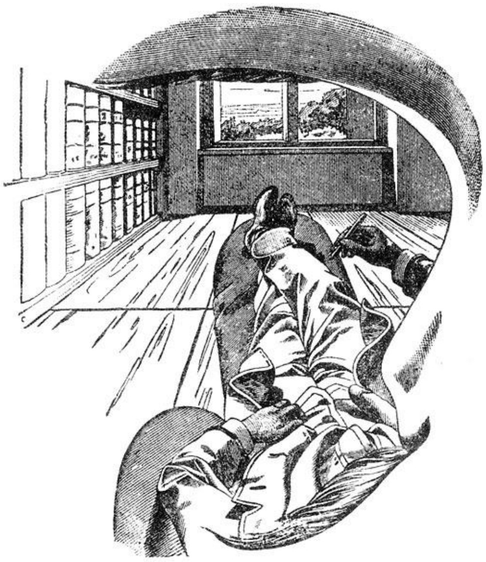

Fig22 . Ernst Mach, Illustration from *The Analysis of Sensations.*

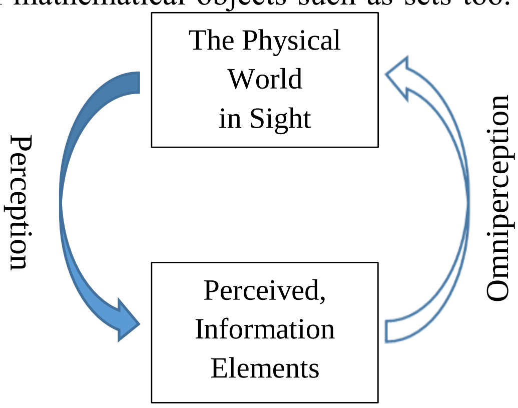

Fig22 . Omniperception.
The full arrows represent the flow of information into consciousness, whereas the hollow arrow represents the experiential leap (reversal).

But it is beside the particular point of this chapter, and so I put it aside and elaborate on the other aspect.

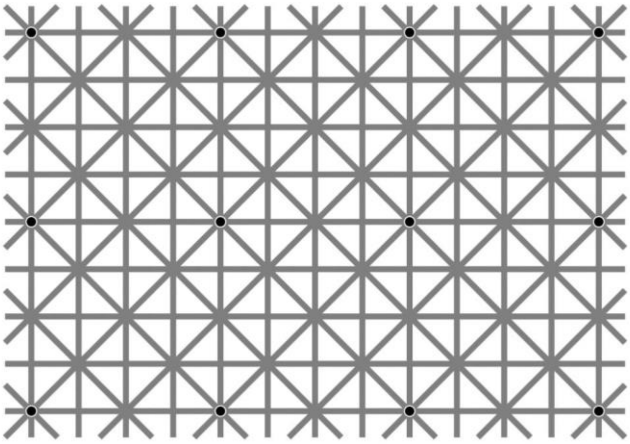

Fig23 . Can you see the twelve black dots all at once? *Ninio's extinction illusion* (adapted by Akiyoshi Kitaoka from *(Ninio, et al., 2000)*).

The other aspect of completeness, which has been alluded to more clearly by Noë and Pinker, is the (alleged) perception of the complete visual scene, i.e., all objects in the scene. Another vivid example (to add to Pinker's) is "Ninio's extinction illusion" (Fig . 23), which, too, may serve to show how, "although a visual scene appears well-defined everywhere, one sees in detail only the portion which is really attended to.. so that a good deal of filling-in may be used without being noticed" (Ninio, et al., 2000).

Peripheral vision serves as a clear albeit simplistic version of this sense in which we seem to perceive the whole that is in front of us and how we may be overestimating it. The relative phenomenological transparency of the motoric activity of the eye itself is brought about by the strong automatic connection between attention and gaze direction. However, the role of attention (which is a very general theme that comprises many different aspects) runs much deeper and is much more general. Attending to particular aspects of the visual scene determines much of the small concrete chunks of information about that scene that we come to consciously perceive. The (large) extent to which attention implicitly mitigates between the particulars and the whole (that is, without our noticing its extensive role) forms a substantial part of the leap.

Fig24 . The Invisible Gorilla / Daniel Simons and Christopher Chabris.

One aspect of attention is *goal-directedness.* We can be extremely focused on very particular aspects of the scene, perceiving just what we need for what we consider the task to be, while the rest is filtered as noise. We thus miss out on what we would otherwise consider central and impossible to miss (such as a large Gorilla taking its time passing through the middle of the scene; Fig . 24).

This may seem as just an extreme, non-representative case, being particularly focused. But the phenomenon of *change-blindness* (Eysenck, et al., 2010 pp. 143-149), “the failure to detect that an object has moved, changed, or disappeared”, uncovers a considerable scope of unaware inawareness. “As a result [of omniperception], we are confident we could immediately detect any change in the visual environment provided it was sufficiently great”. We do have a short-length visual memory buffer that allows us to notice change (as an event in time) if it happens before our very eyes. But it is 'primitive'; the information in it does not necessarily enter our full-blown working memory or consciousness. These change-detection mechanisms are more delicate than we think, heuristic, susceptible in principle to various factors that, for the most part, are simply not part of ordinary life. That we cannot process the *difference* between *separated* scenes (Fig25 .) seems to suggest that we are not in possession of their contents (at a sufficiently high level). Amazingly much of the scene can change without our notice (including even the people we are talking to! (Simons, et al., 1998)).

Fig25 . Tourists / Ronald A. Rensink. Can you spot the difference between the pictures?

We are aware that we do not, strictly speaking, perceive *all* that is perceivable. But it does seem that we perceive it "for the most part" – and for good reason! Pragmatically, probabilistically (– properly weighted), this is indeed correct. Aided by the structure of the world (causal, etc.), heuristic methods of attention behavior suffice, doing the best they can with bounded, finite resources, confronted with an immensely richer reality. The implicit idealization over attention is thus a fundamental component of omniperception. Perceived reality, for us, is everything we *could* attend to. This includes all locations in the scene and, in this way, entails the snapshot conception. It includes events in time (e.g., a moving object). It also includes the non-localized ability to see that it is twilight time; that the car we are in is moving pretty fast; that there are many more people on the other side of the road; etc. However, it does not include, for example, color patterns for colors that we physiologically cannot detect.

8.3.2 Omniperception in the Abstract

Omniperception is most explored, and most obvious, in the field of *visual* perception (though it still lacks a clear theoretical framework even there). This field is quite uniquely rich, and we are (almost) all experts in it. My claim here, however, is that the phenomenon need not be unique to vision; it may be grounded in more fundamental architectural principles, and be far more general. It might be evident in other classical modalities of perception, such as hearing. But in order to see it for what it is in mathematics (if indeed it is there too), it would need to be generalized beyond *sensory* perception altogether; to be understood and put into abstract terms (as I have striven to do). Let us quickly recap and consider this phenomenon in this general context, with *objects*, *reality*, and *perception* no longer restricted *physical* objects and reality and its *visual* perception.

The richer the realm *and* the realm-related cognitive abilities are (with more particular types of perception and combinations of such into higher object-types) – the richer the perceivable (part of) reality can be. (For instance, fewer vs. more types of color cones that reveal more colors that there are to reality). However, a richer perceivable reality channeled into an ultimately limited resource that is working memory, in general determines a larger gap. Some of this can be offset by richer, higher-level object-representations. But not all can be compensated for. Besides possible computational limitations on richness and high-levelness, the number of objects, for example, is still tightly bound. More than a few mathematical objects or work tasks (etc.), too, cannot all be individually perceived at once. This puts the weight on selective attention to make up for it. Attention fixates on finite, very particular aspects of reality, which can be captured in a cognitive representation. These include, first and foremost, the individuation of what we usually come explicitly to take for objects in the realm, such as an apple or an ordinal or a task, perceived into a mental representation of this object-type. Attention-fixation determines a specific domain of relevance, a certain "aspect" of reality (and can do so quite strongly, as in Fig24 .).

A rich role for attention widens the gap, pushing perceivable reality far beyond the actually perceived. However, the machinery seamlessly embeds the actually perceived within an illusory de-cognitivized perception of the full perceivable reality, which "appears well-defined everywhere". Implicit idealization over the focusing role of attention then enlarges the leap, to that perceivable reality as actually perceived.

To sum up: There are many extreme limitations (both practical and principled) tying down whatever is actually *perceived* from some reality into proper representational forms ("the gap"). These limitations do not apply to the holistic, trans-cognitive phenomenological idealization of *perceivable* reality ("the leap"). This is what brings about the illusive impression of omniperception.

8.3.3 Omniperception in Naïve Set Theory

I have set out to explain the intuition that the set-theoretic universe U – is itself is a set. At the core of my explanation, which we now reach, is the analogy that takes this for *a case (in the relevant realm) of omniperception*.

Even if omniperception is indeed a general phenomenon as suggested, positing the perception of sets as on par with the visual perception of physical objects is admittedly a questionable move. Vision is essential and central to our lives as humans, and so we have come to have purpose-specific mechanisms. This means that the non-conscious part that makes up for such a rich area of scientific study may not have an analog in mathematics. I do not defend this move here in general. I hope that my account in this chapter, if taken as reasonable, suggests that there is some validity to it. Just to mention a central historical representative who was willing to go this way, Gödel famously suggested "that we are dealing with [the entities of the set universe] in a manner that is analogous with a natural scientist's study of natural objects. (Godel 1944) In this analogy, a scientist's sense perception has according to Godel a counterpart in mathematical intuition" (Hintikka, 2005). My approach adopts his suggestion only to turn it on its head, as modern cognitive science calls to explore and question the perceptual process itself, rather than simply take perception and its products at face value, as a direct, infallible channel to the realm and its truths.

Fig26 . illustrates the proposed analogy, which I now elaborate.

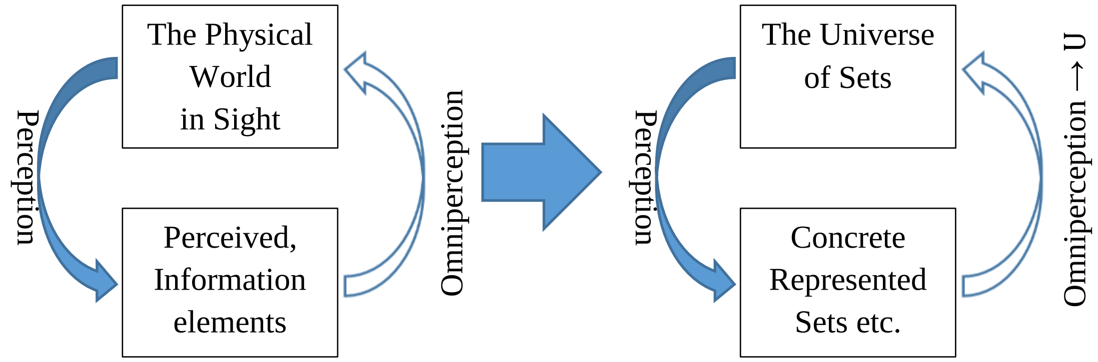

Fig26 . Omniperception in set theory, in analogy from vision (Fig22 .)

On the right: the full arrow represents perception, from the universe of sets, through non-conscious computation, to what in the end is actually perceived. The hollow arrow is the leap: the holistic, realist (de-cognitivized) phenomenology produced over and above the particular bits of perception.

To make the analogy more straightforward for simplicity's sake, let us admit an external, platonic reality of sets, which is somehow being perceived. It is not clear what the cognitive starting point could be, but at least its particulars should, in principle, be inconsequential to the phenomenon of omniperception (as argued on p. 249). What is important is to consider sets in terms of the cognitively, computationally grounded set-*representations* (i.e., along the lines of Section 8.2.2). These elements of perception serve here as the basic content which cognition creates and on which it operates. Importantly, these representations are sufficient for tracing all of the concrete sets (as well as construction principles etc.) that the student has met throughout the basic naïve course. They are of a constructive, bottom-up nature, and their elements are never collected through mere language.

With respect to these representations, within the cognitive system that manages them, all the principles of cognition that have previously been mentioned apply (as is our default GCP assumption, p. 118). So, to continue the portrait begun in Section 8.2.3, of how, through the basic naïve course, the student begins to "master set theory":

She develops expertise to higher and higher levels, at which the individuated computational actions become more and more automated and can continue consciously unattended to (taking a lesser toll on her working memory). Mastered simpler set representations and their operations serve as building blocks (bottom-up) for higher-order ones (just like implicitly finite ordinals and cardinals were amalgamated into numbers, many years earlier; Chapter 7). When no longer needed, the student can then "forget about" lower-level building blocks. Phenomenologically, conscious attention is by default directed at the highest-level objects of relevance.[112]

The richer the set-theoretic cognitive apparatus of the student becomes (with more particular types of set representations, operations on these, and combinations of such into higher types) – the richer her perceivable reality of sets becomes too. (ZFC is one form of retroactively spanning that reality). More kinds of set representations enrich the seemingly a priori (but undefined) general notion of a set. The student "learns to feel comfortable around" more complicated sets, to examine their properties, and "see" how operations are applied to them to produce new ones. Through interaction with the realm (much of which has probably going on much before entering university), she comes to master it.

However, her growing mastery of set theory is still forever bound within the fundamental cognitive limitations. And as the reality of particular sets she actually comes to perceive grows, her idealized *perceivable reality*, her realm of sets, grows much faster.[113] For a central example: once actual infinity and the general notion of a subset are accepted and mastered, she becomes able to perceive and attend to, individuate, various particular subsets of ℕ too – while the idealization brings in much more than that, namely *all of them*. The gap widens, more so than for vision (where everything is ultimately finite).

Idealization over the role and of attention and its limitations entails, in this case now mentioned, that *any specific* subset of ℕ could be attended to, individuated. In general, this idealization, recall, is a fundamental component of *omniperception*. Against the growing gap, the magnitude of omniperception, of this leap "back into reality", grows accordingly.

---

[112] For a different mathematical example, consider how more experienced mathematicians can make do with briefer proof sketches. The many details that would have been required to be made explicit for undergraduates (say), now become a burden on the advanced mathematician's mind, who can see the "gist" of the proof more clearly, at a higher level. (Compare with the reversed hierarchy theory, p. 138).

[113] It should be noted that attention has more ways to navigate in that space too. For example, there may be different definitions or constructions that produce the same set. In general, more place for more complex attentional behavior brings with it place for heuristic "cheats", including ones assimilated underneath consciousness.

We thus arrive at the following suggested picture:

Each actual set ever met and attended to by the student was fully perceived, through some cognitive representation. It was a concrete one, perhaps constructed bottom-up from other pre-existing ones, rooted in actual computational processes. Anything done with that set, any role it ever played, was always within a particular, perceived *context*: Attention was directed at but a few other also-represented entities, representations inhabiting working memory (sets, operations, properties to examine, comparisons, etc.).

However: Coming to develop an expertise, the student's sense of omniperception brings her to "see through this" and *decontextualize*; to generalize what can be perceived over and above the cognitive representations, of the concrete sets actually attended to; to disengage from the representations. Phenomenologically, she makes the leap, seeming to perceive the whole world of sets (if only to the vaguest of degrees, and certainly allowing for turning attention to "new" particular sets). The extreme limitedness of the pieces of information that she comes to grasps about this world is of no consequence. Her mind "completes the picture", as if it is all just in front of her, holistically.

In vision, omniperception does not mean that the whole visible scene is conceived as a proper object too. For set theory, though, its constitutive basis was to consider any perceived collection of sets (or objects, in general), however represented – an object of the realm, namely, a set. Thus, as the whole set-theoretic universe seems to be perceived, "wall-to-wall", it serves as a collection of sets too, and it is consequently considered as simply another set. U emerges directly from set theory's form of omniperception; of the seeming perception of *everything* (that can be perceived).

Aftermath

The realm of sets, qua mathematics, presumably has no paradoxes. Its proper perception, with the concrete sets in particular represented by (or even grounded in) actual processes, would not allow for one. Russell et al. can be considered to have demonstrated that the set universe as a whole is unlike all the concrete sets that the students (and mathematicians at large) ever actually come to meet during the course and represent in their mind and through mathematical language. Even if it is in some sense definite, it in principle does not allow for any finite set representation; it is not *essentially finite* (p. 116). In particular, its linguistic expression is not a valid *set* representation (Section 8.2.2.3).

Essential finitism serves here as a modern cognitive rendition of Cantor's distinction of standard transfinite sets from "absolute infinity": A multitude can be constituted in such a way that the assumption that all of its elements 'are together' leads to a contradiction, so that it is impossible to conceive the multitude as a unity, as 'one finished thing.' Such multitudes I call absolutely infinite or inconsistent multitudes" (Ferreirós, 2007 p. 293). Hence, a standard, valid set can still in *some* sense be conceived as 'one finished thing'. This completed conception I thus naturally take to be the (finite, grasped) representation.

In the footsteps of Cantor's distinction, the resulting picture is in line with the direction that formalized set theories have taken since the discovery of the paradoxes, where sets are "constructed" bottom-up from pre-existing sets and thus adhere to the principle of *the limitation of size*. However, the support for this approach and the limitation of size was chiefly pragmatic ("we get all the mathematics we need, and no new contradictions have appeared so far"). The definite cardinality of the sets that ZF directly commits to may just be a by-product of what is metaphysically central: having concrete cognitive representations. A complete cognitive account of set theory (along the lines I have been sketching) would (hopefully) have pointed to the fault in "the set of all sets" anyway. Even if no one has ever come to hit upon such paradox. Even if there *was no* outright paradox to derive.

Non-Conclusion

I have set out to integrate modern cognitive science into the philosophy of mathematics. This journey, as I see it, has just begun. Each and every chapter does not really conclude, but rather calls for further development within this grand endeavor. To conclude the dissertation, then, let us review it from this point of view:

The whole of Part I was a preliminary setup and consideration of a metaphysical cognitive foundational program for mathematics, replacing phenomenology with cognitive science. In Chapter 1 I have articulated the basic motivation explicitly, within an historical cross-disciplinary context that provides further justification. There is need for a point by point scholarly confrontation with philosophical predecessors such as Kant, Frege, Russell, Hilbert, (and in particular) Brouwer, as more modern ones. And mostly, going forward beyond the broad-terms philosophical manifesto, into the structure of the possible representational and especially cognitive foundations, awaits further, deep conceptual work (on route to an actual full-blown metaphysical solution).

Most philosophers and mathematicians share anti-psychologistic sentiments that sway them away from even considering the metaphysical possibility (and to some extent, by association, the methodological, "representational" possibility). In Chapter 2, I have addressed these various worries, attempting to neutralize them as conclusive objections. In particular, I have argued for the *possibility* of a realist such conception, which could, in principle, address them all (with one, AP10, left as a real threat). But that's just preliminary field-work. What is required is a substantial, robust such account that actually does what I argue could be done.

Alternatively to supporting the realist conception, a cognitive approach could set out to explain *away* (i.e. on cognitive grounds) this conception, this seemingly unique nature of mathematics. Some of the representatives I have covered in Chapter 3 did at least address that. But the task that all covered representatives took upon themselves, to revolutionize the philosophy of mathematics in accordance with their cognitive views, is much more general, far more enormous. They all, I think, could gain considerably from a more careful philosophical examination and refinement. But the general program must expand beyond them, to the vast "non-philosophical" area of mathematical cognition. Much in it may yet *bear* on various philosophical questions, even if the academic field's stated purposes are different.

In Part II, I have developed and pursued an approach that centers on the non-metaphysical midway point towards a cognitive foundation – the mental representation of mathematics. In Chapter 4 I have made explicit my particular approach, which passes through *abstract* formulations of familiar cognitive conceptions and theories, to apply them to mathematical topics. This approach is admittedly risky, and it is necessary to explore its validity and value further. It is an early bird, however, in extending mathematical cognition beyond its current elementary, concrete origin, to mathematics as a whole – as is required for the sake of the philosophy of mathematics at large.

Any such approach would stand on cognitive material (empirical findings, conceptualizations, and theory). For the sake of the cognitively less-informed reader, in Chapter 5 I have introduced in brief a few fundamentals, which were in the background for the chapters to follow. At the other, ideal end, everyone would simply be perfectly versed in all of cognitive science. But more realistically, what would be helpful is a cognitive *toolbox* to build on, of the most common and effective material, that bears the most on particulars mathematical topic as they are explored cognitively.

At the heart of my approach stood the consideration of mathematical objects as (basically) analogous to ordinary, non-mathematical ones, where it comes to what they "mean" for the system, to how it represents them and how it uses those representations. In Chapter 6 I have developed an abstract preliminary "objectification" framework that attends to the "challenge, yet to be met, faced by all developmental proposals: to describe the logical space [loosely speaking] in which learners ever acquire new concepts" (Pietroski, et al., 2008). I have come to see this framework as key to taking further my more general approach, and exploring additional mathematical topics. It is thus essential to substantiate and develop the framework itself. A critical part of this is to elucidate and formalize the concept of object-coherence which is at its heart, and explore the complexities and possible mechanisms of coherentation. Other extensions as described throughout the chapter (e.g. for the interaction of multitudes of objects) are also desirable. Last but not least, the framework's connection to the numerous related topics in the literature should, on both the cognitive side and the formal/A.I.

one, be further substantiated and refined, and for the philosophical side, be drawn to begin with.

In Chapter 7 I have brought my approach to the classic topics of "what numbers are" and what happens beyond the finite. Following a cue from Steiner on the delicate interactions between ordinality and cardinality in the finite, I have used my objectification in order to provide a formal-developmental account of how finite ordinals and cardinals naturally come to be integrated into the natural numbers, in contrast with the infinite. This provides a novel (if, to be sure, high-level) account of the inner-structure of the natural numbers, and exposes an illegitimate influence of our cognitive, statistics-centered cognitive system. This is a substantial direct, concrete contribution to the philosophy of mathematics. But even if accepted as a basically correct conceptual framing, much substance awaits in relating it to the extremely rich cognitive and developmental (even educational) research on the topic.

Finally, in Chapter 8, I've offered a cognitive explanation of the intuition for "Comprehension" and how it can arise from a standard first course on set theory. I've located its source in a 'grand illusion' of visual perception I termed "omniperception". Although a common conception among cognitive scientists, I had to clarify it first so that it could be applied to the sketched objectification-style representational content of naïve set theory (in analogy with vision). Some work is still required for substantiating the analogy in order to make the case. Primarily, omniperception must be elucidated further and put into proper abstract terms. Deep theoretical work is thus required on the part of cognitive science. With a sufficiently-detailed account of the representation of set theory too, we could then turn to examine whether the analogy holds (and ultimately, have science validate it). Most of the work is still ahead, then. And yet, this approach is unlike the paradoxes that happened to be discovered and the more limited axiomatizations that have turn out to suffice. It holds the promise of a guided, *methodological* way to explore such *discrepancies* between the mathematical realm and our conscious perception of it. In this way, it may help to further explore what that reality actually is, ontologically and maybe even metaphysically.

Bibliography

Ahissar Merav [et al.] Reverse hierarchies and sensory learning [Journal] // Phil. Trans. R. Soc. - 2008.

Amalric Marie and Dehaene Stanislas Cortical circuits for mathematical knowledge: evidence for a major subdivision within the brains semantic networks [Journal] // Philosophical Transactions of the Royal Society of London B: Biological Sciences. - [s.l.] : The Royal Society, 2018. - Vol. 373. - ISSN: 0962-8436.

Amalric Marie and Dehaene Stanislas Origins of the brain networks for advanced mathematics in expert mathematicians [Journal] // Proceedings of the National Academy of Sciences. - 2016. - Vol. 113. - pp. 4909-4917.

Aytar Y., Vondrick C. and Torralba A. See, Hear, and Read: Deep Aligned Representations [Journal] // ArXiv e-prints. - 6 2017.

Bailly Francis and Longo Giuseppe Mathematics and the Natural Sciences [Book]. - [s.l.] : Imperial College Press, 2011.

Beeson Michael J. Foundations of Constructive Mathematics [Book]. - [s.l.] : Heidelberg: Springer-Verlag, 1985.

Bell John L. The Axiom of Choice [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2015. - Summer 2015.

Benacerraf Paul and Putnam Hilary Introduction [Book Section] // Philosophy of Mathematics: Selected Readings. - [s.l.] : Cambridge University Press, 1984. - 2.

Benacerraf Paul Mathematical Truth [Journal] // The Journal of Philosophy. - 1973. - Vol. 70. - pp. 661-679.

Benacerraf Paul What Numbers Could not Be [Journal] // The Philosophical Review. - [s.l.] : [Duke University Press, Philosophical Review], 1965. - Vol. 74. - pp. 47-73. - ISSN: 00318108, 15581470.

Berteletti Ilaria, Lucangeli Daniela and Zorzi Marco Representation of numerical and non-numerical order in children [Journal] // Cognition. - 2012. - Vol. 124. - pp. 304-313. - ISSN: 0010-0277.

Bishop Errett and Bridges Douglas Constructive Analysis [Book]. - [s.l.] : Springer-Verlag Berlin Heidelberg, 1985.

Brannon Elizabeth M. The development of ordinal numerical knowledge in infancy [Journal] // Cognition. - 2002. - Vol. 83. - pp. 223-240. - ISSN: 0010-0277.

Bridges D.S. and Richman F. Varieties of constructive mathematics [Book]. - [s.l.] : Cambridge Univ. Press, 1987.

Bridges Douglas and Palmgren Erik Constructive Mathematics [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2016. - Winter 2016.

Bryant Peter and Nuñes Terezinha Childrenś Understanding of Mathematics [Book Section] // The Wiley-Blackwell Handbook of Childhood Cognitive Development. - [s.l.] : Wiley-Blackwell, 2010. - ISBN: 9781444325485.

Chai Wen Jia, Abd Hamid Aini Ismafairus and Abdullah Jafri Malin Working Memory From the Psychological and Neurosciences Perspectives: A Review [Journal] // Frontiers in Psychology. - 2018.

Chalmers David J. A Computational Foundation for the Study of Cognition [Report]. - 1994.

Chirimuuta Explanation in Computational Neuroscience: Causal and Non-causal [Article] // The British Journal for the Philosophy of Science. - 2017. - axw034.

Clarke-Doane Justin What Is the Benacerraf Problem? [Book Section] // Truth, Objects, Infinity: New Perspectives on the Philosophy of Paul Benacerraf / ed. Pataut Fabrice. - Cham : Springer International Publishing, 2016. - ISBN: 978-3-319-45980-6.

Cohen Henri and Lefebvre Claire Handbook of Categorization in Cognitive Science (Second Edition) [Book]. - [s.l.] : Elsevier Science, 2017.

Davidson Kathryn, Eng Kortney and Barner David Does learning to count involve a semantic induction? [Journal] // Cognition. - 2012. - Vol. 123. - pp. 162-173. - ISSN: 0010-0277.

Davis Ernest Mathematics as a metaphor: Review of Where Mathematics Comes From by George Lakoff [Journal] // Journal of Experimental and Theoretical Artificial Intelligence. - [s.l.] : Taylor and Francis Ltd., 2005. - Vol. 17. - pp. 305-315. - ISSN: 0952-813X.

Dehaene Stanislas and Brannon Elizabeth (eds.) Space, Time and Number in the Brain: Searching for the Foundations of Mathematical Thought (Attention and Performance) [Book]. - [s.l.] : Academic Press, 2011.

Dehaene Stanislas The Number Sense: How the Mind Creates Mathematics [Book]. - [s.l.] : Oxford University Press, 1997.

Dehaene Stanislas The Number Sense: How the Mind Creates Mathematics, Revised and Updated Edition [Book]. - [s.l.] : Oxford University Press, 2011.

Dehaene Stanislas, Changeux Jean-Pierre and Naccache Lionel The Global Neuronal Workspace Model of Conscious Access: From Neuronal Architectures to Clinical Applications [Book Section] // Characterizing Consciousness: From Cognition to the Clinic? / ed. Dehaene Stanislas and Christen Yves. - Berlin : Springer Berlin Heidelberg, 2011. - ISBN: 978-3-642-18015-6.

Demopoulos William and Clark Peter The Logicism of Frege, Dedekind, and Russell [Book Section] // The Oxford Handbook of Philosophy of Mathematics and Logic / ed. Shapiro Stewart. - [s.l.] : Oxford University Press, 2005.

Dennett Daniel C. Consciousness Explained [Book]. - [s.l.] : Back Bay Books, 1991.

Dennett Daniel C. Surprise, surprise [Journal] // Behavioral and Brain Sciences. - [s.l.] : Cambridge University Press, 2001. - Vol. 24. - pp. 982–982.

Detlefsen Michael Formalism [Book Section] // The Oxford Handbook of Philosophy of Mathematics and Logic / book auth. Shapiro Stewart. - [s.l.] : Oxford University Press, 2005.

Devlin Keith The Math Gene [Book]. - [s.l.] : Basic Books, 2000.

Dummett Michael Frege: Philosophy of Mathematics [Book]. - [s.l.] : Harvard University Press, 1991.

Eysenck Michael W. and Keane Mark T. Cognitive Psychology: A Student's Handbook, 5th edition [Book]. - [s.l.] : Psychology Press, 2005.

Eysenck Michael W. and Keane Mark T. Cognitive Psychology: A Student's Handbook, 6th Edition [Book]. - [s.l.] : Psychology Press, 2010.

Feferman Solomon Predicativity [Book Section] // The Oxford Handbook of Philosophy of Mathematics and Logic / book auth. Shapiro Stewart. - [s.l.] : Oxford University Press, 2005.

Feferman Solomon What kind of logic is "Independence Friendly" logic? [Book Section] // The Philosophy of Jaakko Hintikka / book auth. Auxier Randall E. and Hahn (eds.) Lewis Edwin. - [s.l.] : Open Court, 2006.

Feigenson Lisa Objects, Sets, and Ensembles [Book Section] // Space, Time and Number in the Brain / ed. Dehaene Stanislas and Brannon Elizabeth M.. - San : Academic Press, 2011. - ISBN: 978-0-12-385948-8.

Feigenson Lisa, Dehaene Stanislas and Spelke Elizabeth Core systems of number [Journal] // Trends in Cognitive Sciences. - 2004. - Vol. 8. - pp. 307-314. - ISSN: 1364-6613.

Ferreirós José Labyrinth of Thought: A History of Set Theory and Its Role in Modern Mathematics (2nd edition) [Book]. - [s.l.] : Birkhäuser Basel, 2007.

Ferreiros Jose The Road to Modern Logic-An Interpretation [Journal] // The Bulletin of Symbolic Logic. - [s.l.] : [Association for Symbolic Logic, Cambridge University Press], 2001. - Vol. 7. - pp. 441-484. - ISSN: 10798986.

Field Hartry Realism, Mathematics & Modality [Book]. - [s.l.] : Blackwell, 1989.

Fodor J. A. and Pylyshyn Z. W. How direct is visual perception?: some reflections on Gibsonś "Ecological Approach́ [Journal] // Cognition. - Netherlands : [s.n.], Apr 1981. - 2 : Vol. 9. - pp. 139-96.

Gaifman Haim Non-Standard Models in a Broader Perspective [Book Section] // Nonstandard Models of Arithmetic and Set Theory / book auth. Enayat Ali and Kossak Roman. - [s.l.] : American Mathematical Soc., 2004.

Goldin Gerald A. Counting on the metaphorical [Journal] // Nature. - #sep# 2001. - Vol. 413. - pp. 18-19. - ISSN: 0028-0836.

Goswami Usha The Wiley-Blackwell Handbook of Childhood Cognitive Development, 2nd Edition [Book]. - [s.l.] : Wiley-Blackwell, 2010.

Gowers Timothy (ed.) The Princeton Companion to Mathematics [Book]. - [s.l.] : Princeton University Press, 2008.

Graham George Behaviorism [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2017. - Spring 2017.

Greenwood John D. Understanding the "cognitive revolution" in psychology [Journal] // Journal of the History of the Behavioral Sciences. - 1999. - 1 : Vol. 35. - pp. 1-22.

Hassabis Demis [et al.] Neuroscience-Inspired Artificial Intelligence [Journal] // Neuron. - [s.l.] : Elsevier. - Vol. 95. - pp. 245-258. - ISSN: 0896-6273.

Hassin Ran R. and Sklar Asael Y. The human Unconscious - A Functional Perspective [Book Section] / book auth. Mind Dual-Process Theories of the Social. - 2014.

Hebart Martin N. and Hesselmann Guido What Visual Information Is Processed in the Human Dorsal Stream? [Journal] // Journal of Neuroscience. - [s.l.] : Society for Neuroscience, 2012. - Vol. 32. - pp. 8107-8109. - ISSN: 0270-6474.

Hendricks Vincent F. and Leitgeb (eds.) Hannes Philosophy of Mathematics: 5 Questions, [Book]. - [s.l.] : New York: Automatic Press / VIP, 2006.

Henrich Joseph, Heine Steven J and Norenzayan Ara The weirdest people in the world? [Journal] // Behavioral and Brain Sciences. - 2010. - 2-3 : Vol. 33. - pp. 61-83.

Hersh Reuben What Is Mathematics, Really? [Book]. - [s.l.] : Oxford Univ. Press., 1997.

Hintikka Jaakko The Principles of Mathematics Revisited [Book]. - [s.l.] : Cambridge University Press, 1996.

Hintikka Jaakko What Platonism? Reflections on the thought of Kurt Godel [Journal] // Revue internationale de philosophie. - 2005. - 4 : Vol. 234. - pp. 535-552.

Hodges Wilfrid A Shorter Model Theory [Book]. - New York, NY, USA : Cambridge University Press, 1997. - ISBN: 0-521-58713-1.

Hohol Mateusz [et al.] Bringing Back the Balance: Domain-General Processes Are Also Important in Numerical Cognition [Journal] // Frontiers in Psychology. - 2017. - Vol. 8. - p. 499. - ISSN: 1664-1078.

Holmes M. Randall Alternative Axiomatic Set Theories [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2017. - Winter 2017.

Hsieh Po-Jang, Colas Jaron T and Kanwisher Nancy Unconscious pop-out: attentional capture by unseen feature singletons only when top-down attention is available. [Journal] // Psychological Science. - 2011. - 9 : Vol. 22. - pp. 1220–1226.

Iemhoff Rosalie Intuitionism in the Philosophy of Mathematics [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2016. - Winter 2016.

Kaiser L. [et al.] One Model To Learn Them All [Journal] // ArXiv e-prints. - 6 2017.

Kinzler Katherine D. and Spelke Elizabeth S. Core systems in human cognition [Journal] // Progress in Brain Research. - 2007. - Vol. 164. - pp. 257-264. - From Action to Cognition. - ISSN: 0079-6123.

Koch Christof [et al.] Neural correlates of consciousness: progress and problems [Journal] // Nat Rev Neurosci. - [s.l.] : Nature Publishing Group, a division of Macmillan Publishers Limited. All Rights Reserved., #may# 2016. - Vol. 17. - pp. 307-321. - ISSN: 1471-003X.

Kreisel G. Two notes on the foundations of set-theory [Journal] // Dialectica. - [s.l.] : Wiley, 1969. - Vol. 23. - pp. 93-114. - ISSN: 00122017, 17468361.

Kusch Martin Psychologism [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2015. - Winter 2015.

Lakoff George and Núñez Rafael E. A reply to Bonnie Gold’s review of “Where Mathematics Comes From: How the Embodied Mind Brings Mathematics Into Being.” [Book]. - [s.l.] : Mathematical Association of America Online., 2001.

Lakoff George and Núñez Rafael Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being [Book]. - [s.l.] : Basic Books, 2000.

Leibovich Tali [et al.] From śense of number΄to śense of magnitude΄- The role of continuous magnitudes in numerical cognition. [Journal] // The Behavioral and brain sciences. - England : [s.n.], Aug 2016. - pp. 1-62.

Levy Arnon What was Hodgkin and Huxley’s Achievement? [Journal] // The British Journal for the Philosophy of Science. - 2014. - Vol. 65. - pp. 469-492.

Linnebo Øystein Platonism in the Philosophy of Mathematics [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2017. - Summer 2017.

Linnebo Øystein The Individuation of the Natural Numbers [Book Section] // New Waves in Philosophy of Mathematics / book auth. Bueno Otávio and Linnebo Øystein (eds). - [s.l.] : Palgrave Macmillan UK, 2009.

Longo Giuseppe and Viarouge Arnaud Mathematical intuition and the cognitive roots of mathematical concepts [Journal] // Topoi. - 2010. - 1 : Vol. 29. - pp. 15-27.

Longo Giuseppe Mathematical Concepts and Physical Objects [Book Section] // Rediscovering Phenomenology / book auth. Boi L., Kerszberg P. and Patras (eds.) F.. - [s.l.] : Springer, 2007.

Longo Giuseppe Mathematical Intelligence, Infinity and Machines: beyond the Goedelitis [Journal] // Journal of Consciousness Studies. - 1999. - 11-12 : Vol. 6.

Longo Giuseppe The Cognitive Foundations of Mathematics: human gestures in proofs and mathematical incompleteness of formalisms [Book Section] // Images and reasoning / book auth. Grialou Pierre, Longo Giuseppe and Mitsuhiro Okada¸. - [s.l.] : Keio University Press, 2005.

Longo Giuseppe The Constructed Objectivity of Mathematics and the Cognitive Subject [Book Section] // Quantum Mechanics, Mathematics, Cognition and Action - Proposals in Epistemology. / book auth. Mugur-Schachter (ed.) M.. - [s.l.] : Kluwer, 2002.

Longo Giuseppe The Reasonable Effectiveness of Mathematics and its Cognitive Roots. [Book Section] // Geometries of Nature, Living Systems and Human Cognition - series in "New Interactions of Mathematics with Natural Sciences and Humanities" / book auth. Boi (ed.) L.. - [s.l.] : World Scientific, 2005.

Longo Giuseppe Theorems as Constructive Visions [Conference] // Proceedings of ICMI 19 conference on Proof and Proving, Taipei, Taiwan, May 10 - 15, 2009. - [s.l.] : Springer, 2010.

Lyubova Natalia, Ivaldi Serena and Filliat David From Passive to Interactive Object Learning and Recognition Through Self-identification on a Humanoid Robot [Journal] // Auton. Robots. - Hingham : Kluwer Academic Publishers, #jan# 2016. - Vol. 40. - pp. 33-57. - ISSN: 0929-5593.

Madden James Where Mathematics Comes From - Book Review [Journal] // Notices of the AMS. - November 2001. - 10 : Vol. 48.

Maddy Penelope Set-theoretic Foundations [Book Section] // Foundations of Mathematics / book auth. Caicedo Andres Eduardo [et al.]. - [s.l.] : American Mathematical Society, 2017.

Maddy Penelope Three Forms of Naturalism [Book Section] // The Oxford Handbook of Philosophy of Mathematics and Logic / book auth. Shapiro Stewart. - [s.l.] : Oxford University Press, 2005.

Manin Yuri Mathematics as profession and vocation [Book Section] // Mathematics: Frontiers and Perspectives / book auth. Arnol′d Vladimir Igorevich et. al. - [s.l.] : AMS, 2000.

Miller George A. The cognitive revolution: a historical perspective [Journal] // Trends in Cognitive Sciences. - [s.l.] : Elsevier. - Vol. 7. - pp. 141-144. - ISSN: 1364-6613.

Modayil Joseph and Kuipers Benjamin The initial development of object knowledge by a learning robot [Journal] // Robotics and Autonomous Systems. - 2008. - 11 : Vol. 56.

Moors Agnes and Houwer Jan De Automaticity: a theoretical and conceptual analysis. [Journal] // Psychological bulletin. - 2006. - Vol. 132 2. - pp. 297-326.

Moschovakis Yiannis N. Notes on Set Theory [Book]. - [s.l.] : Springer New York, 1994.

Nanay Bence Between Perception and Action [Book]. - [s.l.] : Oxford University Press, 2013.

Ninio Jacques and Stevens Kent A. Variations on the Hermann Grid: An Extinction Illusion [Journal] // Perception. - 2000. - Vol. 29. - pp. 1209-1217. - PMID: 11220212.

Noë Alva Is the Visual World a Grand Illusion? [Book]. - [s.l.] : Imprint Academic, 2002.

Núñez Rafael E. Is There Really an Evolved Capacity for Number? [Journal] // Trends in Cognitive Sciences. - [s.l.] : Elsevier, 6 2017. - 6 : Vol. 21. - pp. 409-424. - ISSN: 1364-6613.

O'Connor Patrick A., Morsanyi Kinga and McCormack Teresa Young children's non- numerical ordering ability at the start of formal education longitudinally predicts their symbolic number skills and academic achievement in maths [Journal] // Developmental Science. - 2018.

Ohshiro Tomokazu, Angelaki Dora E. and DeAngelis Gregory C. A normalization model of multisensory integration. [Journal] // Nature neuroscience. - United : [s.n.], Jun 2011. - 6 : Vol. 14. - pp. 775-82.

Ojose Bobby Applying Piaget's Theory of Cognitive Development to Mathematics Instruction [Journal] // The Mathematics Educator. - 2008. - Vol. 18.

Paseau Alexander Naturalism in the Philosophy of Mathematics [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2016. - Winter 2016.

Pease Alison, Guhe Markus and Smaill Alan Special Issue: Developments in Research on Mathematical Practice and Cognition [Journal] // Topics in Cognitive Science. - April 2013. - 2 : Vol. 5.

Piccinini Gualtiero Computationalism [Book Section] // The Oxford Handbook of Philosophy and Cognitive Science / book auth. Margolis Eric, Samuels Richard and Stich Stephen P.. - 2012.

Pietroski Paul and Lidz Jeffrey Natural number concepts: No derivation without formalization [Journal] // Behavioral and Brain Sciences. - [s.l.] : Cambridge University Press, 2008. - Vol. 31. - pp. 666-667.

Pinker Steven The Brain: The Mystery of Consciousness [Article] // Time Magazine. - 2007.

Pitt David Mental Representation [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2017. - Spring 2017.

Posy Carl Computability and Constructability [Book Section] // Computability: Turing, Gödel, Church, and Beyond / book auth. Copeland Posy, Shagrir (eds). - [s.l.] : MIT Press, 2013.

Posy Carl Intuitionism [?] [Book]. - forthcoming.

Posy Carl Intuitionism and Philosophy [Book Section] // The Oxford Handbook of Philosophy of Mathematics and Logic / ed. Shapiro Stewart. - [s.l.] : Oxford University Press, 2005.

Posy Carl Of Griffins and Horse: Mathematics and Metaphysics in Kant’s Critical Turn [Book Section] // Kant’s Philosophy of Mathematics, Volume 1: The Critical

Philosophy and its Roots / book auth. Posy Carl and Rechter Ofra (eds.). - [s.l.] : Cambridge University Press, Forthcoming, 2019.

Prakash Ravi [et al.] Global workspace model of consciousness and its electromagnetic correlates [Journal] // Annals of Indian Academy of Neurology. - 7 2008. - 3 : Vol. 11. - pp. 146-53.

Putnam Hilary Mathematics without Foundations [Journal] // The Journal of Philosophy. - [s.l.] : Journal of Philosophy, Inc., 1967. - Vol. 64. - pp. 5-22. - ISSN: 0022362X.

Quinn Paul C. Born to Categorize [Book Section] // The Wiley-Blackwell Handbook of Childhood Cognitive Development, 2nd Edition / book auth. Goswami Usha. - [s.l.] : Wiley-Blackwell, 2010.

Rayo Agustín and Uzquiano Gabriel Absolute Generality [Book]. - [s.l.] : Oxford University Press, 2006.

Rescorla Michael The Computational Theory of Mind [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2017. - Spring 2017.

Rips Lance J. Beliefs about the nature of numbers [Book Section] // Mathematics, Substance and Surmise. - [s.l.] : Springer International Publishing, 2015. - ISBN: 9783319214726.

Rips Lance J., Bloomfield Amber and Asmuth Jennifer From numerical concepts to concepts of number [Journal] // Behavioral and Brain Sciences. - [s.l.] : Cambridge University Press, 2008. - Vol. 31. - pp. 623-642.

Robinson J. O. The psychology of visual illusion [Book]. - [s.l.] : Hutchinson London, 1972. - pp. 288, [3] p. . - ISBN: 0091122805 0091122813 .

Russell Bertrand Introduction to Mathematical Philosophy [Book]. - [s.l.] : Dover Publications, 1919.

Samuels Richard, Shapiro Stewart and Snyder Eric Basically Natural: On the Relative Fundamentality of Ordinal and Cardinal Conceptions of Number [Conference] // Forthcoming?. - 2017.

Schiralli Martin and Sinclair Nathalie A Constructive Response to "Where Mathematics Comes From." [Journal] // Educational Studies in Mathematics. - 2003. - Vol. 52. - p. 79. - ISSN: 00131954.

Schneider Walter and Chein Jason M. Controlled & automatic processing: behavior, theory, and biological mechanisms [Journal] // Cognitive Science. - 2003. - Vol. 27. - pp. 525–559.

Shabel Lisa Kant's Philosophy of Mathematics [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2016. - Spring 2016.

Shapiro Stewart and Wright Crispin All Things Indefinitely Extensible [Book Section] // Absolute Generality / book auth. Rayo Agustín and Uzquiano Gabriel. - [s.l.] : Oxford University Press, 2006.

Shapiro Stewart Foundations of Mathematics: Metaphysics, Epistemology, Structure [Journal] // The Philosophical Quarterly. - 2004. - 214 : Vol. 54. - pp. 16–37.

Shapiro Stewart Foundations without Foundationalism [Book]. - [s.l.] : Oxford University Press, 2000.

Shapiro Stewart Philosophy of Mathematics: Structure and Ontology [Book]. - [s.l.] : Oxford University Press, 1997.

Shapiro Stewart We hold these truths to be self evident: But what do we mean by that? [Journal] // The Review of Symbolic Logic. - [s.l.] : Cambridge University Press, 3 2009. - Vol. 2. - pp. 175-207. - ISSN: 1755-0203.

Simons Daniel J. and Levin Daniel T. Failure to detect changes to people during a real-world interaction [Journal] // Psychonomic Bulletin & Review. - 12 01, 1998. - Vol. 5. - pp. 644-649. - ISSN: 1531-5320.

Sklar Asael Y. [et al.] Reading and doing arithmetic nonconsciously [Journal] // Proceedings of the National Academy of Sciences. - November 27, 2012. - 48 : Vol. 109.

Sloman Aaron Construction Kits for Biological Evolution [Book Section] // The Incomputable: Journeys Beyond the Turing Barrier / ed. Cooper S. Barry and Soskova

Mariya I.. - Cham : Springer International Publishing, 2017. - ISBN: 978-3-319-43669-2.

Sloman Aaron Diagrams in the Mind? [Book Section] // Diagrammatic Representation and Reasoning / book auth. Anderson Michael, Meyer Bernd and Olivier Patrick. - [s.l.] : Springer London, 2002.

Sloman Aaron Kantian Philosophy of Mathematics and Young Robots [Conference] // {Intelligent Computer Mathematics} / ed. Autexier S. [et al.]. - Berlin/Heidelberg : Springer, 2008. - pp. 558-573.

Sloman Aaron Natural Vision and Mathematics: Seeing Impossibilities [Conference] // Second Workshop on: Bridging the Gap between Human and Automated Reasoning, at IJCAI 2016 / ed. Furbach Ulrich and Schon Claudia. - New : [s.n.], 2016. - pp. 86-101. - http://ceur-ws.org/Vol-1651/.

Sloman Aaron The Computer Revolution in Philosophy: Philosophy, science and models of mind [Book]. - [s.l.] : Harvester press and Humanities press, 1978.

Sloman Aaron The Well-Designed Young Mathematician [Journal] // Artificial Intelligence. - [s.l.] : Elsevier, 2008. - Vol. 172. - pp. 2015-2034.

Snyder Eric, Samuels Richard and Shapiro Stewart Neologicism, Frege's Constraint, and the Frege-Heck Condition [Journal] // Noûs. - 2018. - pp. 1-24.

Snyder Eric, Shapiro Stewart and Samuels Richard Cardinals, Ordinals, and the Prospects for a Fregean Foundation [Journal] // Royal Institute of Philosophy Supplement. - [s.l.] : Cambridge University Press, 2018. - Vol. 82. - pp. 77-107.

Stein Barry E., Stanford Terrence R. and Rowland Benjamin A. Development of multisensory integration from the perspective of the individual neuron [Journal] // Nat Rev Neurosci. - [s.l.] : Nature Publishing Group, a division of Macmillan Publishers Limited. All Rights Reserved., #aug# 2014. - Vol. 15. - pp. 520-535. - ISSN: 1471-003X.

Steiner Mark Mathematics — Application and Applicability [Book Section] // The Oxford Handbook of Philosophy of Mathematics and Logic / book auth. Shapiro Stewart (ed.). - 2005.

Sury Dana and Rubinsten Orly Ordinal Processing of Numerical and Non-numerical Information [Book Section] // Reading, Writing, Mathematics and the Developing Brain: Listening to Many Voices / ed. Breznitz Zvia [et al.]. - Dordrecht : Springer Netherlands, 2012. - ISBN: 978-94-007-4086-0.

Tait W.W. Frege versus Cantor and Dedekind: On the Concept of Number [Book Section] // Frege: Importance and Legacy / book auth. Schirn M.. - [s.l.] : Berlin: de Gruyter, 1996.

Ugur Emre, Oztop Erhan and Sahin Erol Goal Emulation and Planning in Perceptual Space Using Learned Affordances [Journal] // Robot. Auton. Syst.. - Amsterdam, The Netherlands, The Netherlands : North-Holland Publishing Co., #jul# 2011. - Vol. 59. - pp. 580-595. - ISSN: 0921-8890.

van den Berg Ronald, Awh Edward and Ma Wei Ji Factorial comparison of working memory models [Journal] // Psychological Review. - 2014. - 1 : Vol. 121. - pp. 124-149.

van Hoof H. [et al.] Maximally informative interaction learning for scene exploration [Conference] // 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. - 2012. - pp. 5152-5158. - ISSN: 2153-0858.

Voorhees Burton Embodied Mathematics [Journal] // Journal of Consciousness Studies. - 2004. - 9 : Vol. 11. - pp. 83–8.

Wagner Roi Making and Breaking Mathematical Sense: Histories and Philosophies of Mathematical Practice [Book]. - [s.l.] : Princeton University Press, 2017.

Wiki-Afford Affordance --- Wikipedia, The Free Encyclopedia. - 2018. - [Online; accessed 23-October-2018].

Wiki-APiori A priori and a posteriori --- Wikipedia, The Free Encyclopedia. - 2018. - [Online; accessed 23-October-2018].

Wiki-EC Embodied cognition --- Wikipedia, The Free Encyclopedia. - 2018. - [Online; accessed 23-October-2018].

Wiki-ISL Inverse-square law --- Wikipedia, The Free Encyclopedia. - 2018. - [Online; accessed 23-October-2018].

Wiki-MSInt Multisensory integration --- Wikipedia, The Free Encyclopedia. - 2018. - [Online; accessed 23-October-2018].

Wiki-PoG Principles of grouping --- Wikipedia, The Free Encyclopedia. - 2018. - [Online; accessed 23-October-2018].

Wikip-Ord Ordinal number --- Wikipedia, The Free Encyclopedia. - 2018. - [Online; accessed 23-October-2018].

Wilson Robert A. and Foglia Lucia Embodied Cognition [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2017. - Spring 2017.

Zach Richard Hilbert's Program [Book Section] // The Stanford Encyclopedia of Philosophy / ed. Zalta Edward N.. - [s.l.] : Metaphysics Research Lab, Stanford University, 2016. - Spring 2016.

Zach Richard Hilbert's Program Then and Now [Book Section] // Philosophy of Logic / ed. Jacquette Dale. - Amsterdam : North-Holland, 2007. - ISSN: 18789846.

Zalta Edward N. Frege's Theorem and Foundations for Arithmetic [Book Section] // The Stanford Encyclopedia of Philosophy. - [s.l.] : Metaphysics Research Lab, Stanford University, 2018. - Fall 2018.

Zalta Edward N. Gottlob Frege [Book Section] // The Stanford Encyclopedia of Philosophy. - [s.l.] : Metaphysics Research Lab, Stanford University, 2017. - Spring 2017.