UNIVERSITÉ DU QUÉBEC À MONTRÉAL

NUMBERS BEFORE NUMERALS: THE LIMITS OF EXTERNALIST ACCOUNTS OF NUMERICAL COGNITION

THESIS
SUBMITTED TO UQAM
IN PARTIAL FULLFILLMENT OF THE REQUIREMENTS
OF THE DEGREE OF PHD IN PHILOSOPHY

BY
JEAN-CHARLES PELLAND

JANUARY 2019

UNIVERSITÉ DU QUÉBEC À MONTRÉAL

LES NUMÉRAUX AVANT LES NOMBRES ?
LES LIMITES DES MODÈLES EXTERNALISTES
DE L'ORIGINE DES NOMBRES

THÈSE
PRÉSENTÉE
COMME EXIGENCE PARTIELLE
DU DOCTORAT EN PHILOSOPHIE

PAR
JEAN-CHARLES PELLAND

JANVIER 2019

ACKNOWLEDGEMENTS

While I have generally been sceptical of the worth of presenting acknowledgements at the beginning of a book, given that most readers do not know or care about the people the author wishes to thank, and most people mentioned in the acknowledgements have nothing to do with the material presented in the main text, I have come to appreciate the value of publicly thanking those who have contributed to the realization of projects such as this doctoral dissertation. For while it is a potentially uninteresting preface to the work for the anonymous reader, for the author it represents perhaps the only opportunity to express his or her thanks to those who have helped the project become a reality. Also, there is a sense in which putting such thanks in writing adds more weight to them, given the relative permanence of written over oral communication. Plus, in cases such as this dissertation, the acknowledgements might be the only part of the text that those who have helped it along the way can understand! With these thoughts in mind, then, I will allow myself a few thanks of my own, in no particular order, with the the following explicit warning: these thanks in no way guarantee the present work's worth and in no way are those being thanked in any way responsible for its inevitable shortcomings. In short, for better or for worse, I take full responsibility for what lies below.

First I would like to thank my director, Mathieu Marion. I suppose most of all I wish to thank him for his confidence: over the many years I have known him – first during my masters and then my PhD – Mathieu has managed to keep to himself any doubts he might have had about my abilities to carry out my projects. He has always let me do things my way and never asked me to back down from my ideas, no matter how ridiculous many of them turned out to be. I have been lucky to have him watch over me from a distance over these years. It has always been a comfort to me to know that, should I chose to, I could have access to a brilliant mind whose knowledge of foundations of mathematics still amazes me. I have also been blessed with his friendship. Thanks doc M.

I would also like to thank other professors and fellow students at the philosophy department of Université du Québec à Montréal for their support. In particular, Luc Faucher has been a valuable asset. His thorough and detailed reading and criticism of my work have considerably improved it and helped me significantly broaden the scope of my research, on top of injecting it with a much-needed reality check. I also want to thank Alain Voizard for his help and wisdom over the years, even though he is responsible for introducing me to Kripkenstein, which was both a curse to my thesis and a mind-warping philosophical experience. I would be remiss if I did not thank Mauro Rossi for letting me feed off his experience and for sharing his invaluable advice throughout the years, as well as for his unusually pleasant company. Of course, some of my fellow students often provided much needed moral support throughout the long years during which I worked on this dissertation, including Sarah "Ssssarrrahhhh" Arnaud, Denis "Destination Chernobyl" Courville, and Alice "lvdrlvdrlvdr" Livadaru. Merci mes ami-e-s.

I would also like to take this opportunity to thank those whom I met during my time spent at the institute of philosophy in London. In particular, Ophelia Deroy has significantly affected my thinking on many issues presented here. Merely hanging around the IP in London allowed me to soak in great ideas and radically expanded my understanding of how to factor in empirical data to study philosophical issues. I was lucky to benefit from many great conversations with Ophelia and people like Nick Shea, as well as with my dear friend Laura Silva. Being in London also allowed me to spend time with Brian Ball, who was kind enough to read and comment on parts of this thesis and engage in long discussions on a variety of topics of philosophical interest, many of which ended up in this thesis. I treasure our friendship and am lucky to be able to benefit from Brian's expertise. Thanks B.

Thanks also to Snoo for her support in London during this part of my life. So much love.

I have been blessed to be surrounded by a family that has always supported me. My parents have given me not only material and financial support, but also unwavering moral support, even if it was sometimes in the form of guilt trips. They are both sources of inspiration for me, in completely different ways, and I am thankful to have had them as models to look up to when my confidence in this project faltered. It's tough to be a cynic when people are actually accomplishing things all around you. My sister was also a source of inspiration for me, as well as a source of counsel and support, both emotional and physical. I thank her for sacrificing her time at the cabin to allow me to reap the benefits that solitude in the woods offered me, and, of course, the use of her adjustable table! Her patience and understanding in the face of my deteriorating mental state throughout the gruelling last months of writing this thesis should not go unmentioned. I know no one who is as close to their family as I am. Merci beaucoup, je vous aime.

TABLE OF CONTENTS

LIST OF FIGURES

Les numéraux avant les nombres ?
Les limites des modèles externalistes de l'origine des nombres

RÉSUMÉ

Propulsée par le développement de nouvelles technologies et méthodes d'investigation dans divers domaines de recherche, l'étude de la cognition numérique progresse à un rythme fulgurant depuis quelques années. Des domaines aussi variés que la psychologie développementale, l'anthropologie, la linguistique, la neuropsychologie, l'éthologie, et la philosophie contribuent tous à une explosion de données qui nous permettent de croire que l'on pourrait bientôt percer le mystère de comment notre système nerveux pourrait nous permettre de représenter des entités objectives comme les nombres naturels. Une des plus importantes découvertes tirées de ce progrès est celle de ce qu'il est maintenant coutume d'appeler le « Approximate Number System » (ANS), un ensemble de neurones qui nous permet de déterminer la quantité approximative d'objets dans des collections auxquelles nous portons notre attention. Un autre système, le « Object-File System » (OFS), serait quant à lui dédié à garder en tête les propriétés spatiotemporelles d'un nombre restreint d'objets. Malheureusement, les limitations évidentes de ces systèmes indiquent qu'ils ne peuvent à eux seuls expliquer l'émergence de nos capacités arithmétiques : tandis que le ANS est sévèrement limité dans sa précision, la portée du OFS est limitée à un maximum de 4 objets. Considérant ces limitations, nous sommes donc encore loin de savoir comment nous obtenons des représentations avec la précision et l'objectivité des concepts de nombres utilisés dans les mathématiques à partir de systèmes comme le ANS et le OFS.

Pour expliquer comment nous arrivons à traverser le fossé conceptuel entre le contenu quantitatif produit par nos cerveaux et celui qui nous permet de pratiquer l'arithmétique formel, plusieurs modèles de l'origine de nos concepts de nombres ont été offerts récemment. Dans ma thèse de doctorat, je propose une analyse critique des modèles dominants du développement historique et ontogénétique des concepts de nombre. Après un survol des principales données empiriques dans les deux premiers chapitres, je présente un résumé des grandes lignes de deux théories souvent citées dans la littérature sur la cognition numérique, soit celle du sens des nombres de Stanislas

Dehaene (2011) dans le chapitre trois et celle de la cognition de base de Susan Carey (2009) dans le chapitre quatre. Bien qu'elles fassent appel à des processus bien distincts, ces théories se fient à l'existence de symboles numériques dans l'environnement pour expliquer le développement de nos concepts de nombre. L'idée ici serait que notre esprit et les symboles numériques forment un seul système cognitif dans lequel tant notre cerveau que les symboles jouent un rôle actif dans le développement des concepts de nombre, en conformité avec l'externalisme actif proposé par Andy Clark et David Chalmers dans leur théorie de l'esprit étendu (1998). Or, selon l'argument principal de ma thèse, l'existence de ces symboles numériques présuppose la présence des mêmes concepts de nombre dont on tente d'expliquer l'origine. En se fiant à une interaction entre des symboles numériques et notre esprit pour expliquer l'origine de nos concepts de nombre, on place donc la charrue devant les bœufs.

Pour démontrer que cette position est problématique, j'analyse l'approche externaliste à la cognition numérique de Catarina Dutilh Novaes (2013) et de Lambros Malafouris (2010) dans le chapitre cinq pour déterminer à quel point ces théories sont en mesure d'expliquer le développement de contenu conceptuel numérique dans un environnement qui ne contient pas de symboles pour des nombres. Je prétends que ces théories sont incapables de préciser ce qui différencie un individu ayant développé des concepts de nombre d'un individu sans concepts de nombre quand ces individus ont tous les deux accès au même environnement. Par la suite, j'évalue à quel point l'évolution culturelle ou l'extension de la cognition pour inclure notre environnement culturel peuvent aider l'approche externaliste dans le chapitre six. Je présente des arguments selon lesquels ces options ne peuvent aider l'externaliste, puisqu'ils font appel à des mécanismes actifs à l'échelle d'une population, tandis que la cognition numérique a lieu au niveau de l'individu. Si mes arguments tiennent la route, l'externalisme dans l'étude des fondements de la cognition numérique ne permet pas d'expliquer le développement de concepts de nombre dans un environnement sans soutien externe à cette cognition. L'approche externaliste à la cognition numérique devrait donc être limitée aux cas qui requièrent des soutiens externes au-delà d'un segment initial des nombres naturels, dont l'origine doit être expliquée en faisant appel à des modèles internalistes.

MOTS-CLÉS : cognition numérique ; esprit étendu ; sens des nombres ; fondements des mathématiques ; évolution culturelle ; cognition de base

ABSTRACT

The study of the cognitive and perceptual systems underlying our numerical abilities has progressed tremendously in the past few decades, yielding scores of data on the potential role played by the so-called Approximate Number System (ANS) and the Object-File System (OFS) in the development of natural number concepts (Dehaene 1997/2011). While there is still disagreement on the relationship between these systems and on the extent to which they produce representations with numerical content, there is overwhelming consensus that, on their own, neither of these systems produces representations with sufficient precision and numerical range to account for the development of natural number concepts.

The question I am interested in in my doctoral thesis is, given these systems' limitations, how do we manage to build representations with mathematically-viable numerical content? That is, how do we bridge the gap between the quantity-related content produced by our evolutionarily ancient brains and the mathematically-viable numerical content associated with numeration systems like Indo-Arabic or Roman numerals? This is what I refer to as the *gap problem*. It is the main problem that concerns us in this thesis. While many answers have been proposed to this question over the years, virtually all of them rely on culturally-inherited symbols and external artefacts to bridge the gap between natural numbers and the output of our innate cognitive machinery. And yet, as I argue in my thesis, if we want to bridge the gap between natural number concepts and the content produced by systems like the ANS and the OFS, such externalist approaches to cognition are limited in their explanatory power. The main problem with externalist accounts is what I call the *origins problem*: how can we explain the origins of numerical cognition by appealing to external symbols for numbers, when these symbols in turn depend for internal representations of number for their origins?

To support my claim that externalism is unable to answer this problem, the first two chapters start by summarizing data concerning the main cognitive systems involved in numerical cognition. Then I present the main lines of two of the most influential externalist accounts, Stanislas Dehaene's *Number Sense* (chapter 3) and Susan Carey's Quinian Bootstrap (chapter 4), in order to illustrate externalist approaches to the gap problem. I show that despite their differences, both accounts attribute a central role to external representations of numbers like number words in explaining how we bridge the gap, and that this seems problematic given that there are historical cases of development of numerical content without such external support. In chapter 5, I take a closer look at the philosophical motivations behind this externalism in Clark & Chalmer's classic 1998 paper on the extended mind before exploring the constitutivity of external supports for cognition, as framed by Catarina Dutilh Novaes (2013). This leads me to discuss the relationship between external and internal representations for numbers at both ontogenetic and historical timescales. To ground this discussion, I then discuss data concerning possible origins of the first abstract symbols for numbers in Sumeria (Malafouris 2010) and then explore the limits of anumerate cultures like the Pirahã and the Munduruku in order to determine whether there is evidence that we can explain the origins of the concept of precise quantity by appealing to the ability to put objects into one-to-one correspondence. I argue that the data do not support this conclusion, and thus that externalists do not have an account of how the first numerical content emerged in a numeral-free environment. In chapter 6, I explore two potential externalist replies to my gap problem based on the potential constructive role of culture. The first is to try to explain the emergence of novel numerical content by appealing to mechanisms of cultural evolution, as described by Helen De Cruz (2007). I argue that this doesn't help, since such mechanisms are population-level, while the generation of novel content occurs at the level of the individual. I then consider Richard Menary's (2015a) enculturated approach to numerical cognition to see if extending cognition to include our cultural niche can help the externalist, and find this option wanting as well. Here, I argue that innovation matters in how we want to answer the gap problem and that enculturation is not well equipped to describe the individual-level construction of novel content, due to its focus on population-level processes of innovation. The upshot is that externalist approaches need to be restricted in order to make room for internalist theories of the development of numerical content for an initial segment of the natural numbers.

KEYWORDS: Numerical cognition; Extended Mind; Approximate Number Sense; Object-File System; Cultural Evolution; Foundations of Mathematics

INTRODUCTION

1 Why foundations of mathematics needs psychology

While the study of numerical cognition was limited to a few brave psychologists until the end of the 20th century, the last twenty years have witnessed a complete reversal of this trend. Nowadays, fields as diverse as developmental psychology, neuroscience, ethology, anthropology, archeology, and philosophy have all contributed to a recent surge of interest into the cognitive foundations of our ability to think about – and with – numbers. While the opening of these experimental floodgates is certainly a step forward into our understanding of the cognitive and neural systems involved in representing numbers, it also poses difficult questions concerning the relation between psychology and mathematics. After all, empirical research into numerical cognition does more than welcome psychological description of mathematical entities: by its very nature, it tries to ground the concept of number in broadly psychological facts.

For a while, identifying or reducing logical and mathematical concepts to the mental processes or representations responsible for them was an acceptable, if not 'prevalent' (Dummett 2001: xxi) approach to logic. Mill, for example, wrote that logic is "not a Science distinct from, and co-ordinate with Psychology. So far as it is a science at all, it is a part, or branch, of Psychology" (Mill 1979, 359). Mill was far from being alone in espousing a form of psychologism with respect to logic and mathematics:[1] in the

[1] Though there is debate as to whether Mill's philosophy of mathematics was in fact a psychologistic, it is clear that some of his writings are sympathetic to this view. See Godden 2005 for a discussion.

years following Hegel's death, many German philosophers (e.g. Wundt, Erdmann, Lipps, Sigwart, Jerusalem, and others)[2] took a naturalist turn in which mathematics and logic were founded in psychology.

However, in addition to this naturalist turn, the 19th century also witnessed the birth of a number of revolutionary developments in mathematics, including group theory, real analysis, as well as non-Euclidean and projective geometries, to name a few. With these new mathematical practices came increasingly abstract concepts, which spurred the need to secure solid conceptual foundations for mathematics, to ensure that such innovations were free from paradox, contradiction, and the subjectivity of human intuition. It is of little surprise, then, that these innovations in mathematical practice generally steered clear of the relationship between mind and mathematics. After all, if one of the main foundationalist goals is to guarantee the generality and objectivity of mathematics, then accepting concerns related to how our subjective mind interacts with mathematical entities might seem counterproductive. At best, such empirical concerns would be irrelevant. Arguably, these foundational efforts – and the accompanying anti-psychologism – reached their apex with Frege's definition of number in his *Grundlagen* (1884). There, Frege defined number using only logical tools, thus laying the conceptual groundwork needed to found mathematics on the apparently objective and paradox-free edifice of logic.

Frege's anti-psychologistic comments in his *Grundlagen* and, to a lesser extent, in the foreword to his *Grundgesetze* (1893) contained many of the same arguments that would be developed in a more detailed and systematic attack in Husserl's *Logical investigations* (1900).[3] Though discussing these arguments and their similarities is not

---

[2] See Kusch 1995 for a more exhaustive list.

[3] It is worth mentioning that a likely (though contested) reason Husserl renounced the psychologism of his *Philosophy of Arithmetic* (1891) was Frege's (1894) 'savage' critique of this work (Dummett 2001, xxi). See also Mohanti 1982 and Drummond 1985 for more on the relationship between their views.

necessary here, it is worth observing that both Frege and Husserl held that any appeal to psychological notions in explaining the validity of logical or mathematical laws led to some form of relativism, subjectivism, or idealism – all of which run counter to the objectivity, precision, and universality of mathematical truth.[4] On their view, mathematical facts concern eternal and objective platonic entities, whereas psychological facts are about subjective, relative, and temporal things.

Thus, in stark contrast to Mill, Frege wrote that "number is no whit more an object of psychology or the product of mental operations than, let us say, the North Sea is" (Frege 1884/1960, 34). Further, were we to accept the relevance of the psychological origins of an idea when defining or proving it, Frege claims we would be forced to include psychological facts in mathematical proofs: "in proving Pythagoras' theorem we should be reduced to allowing for the phosphorous content of the human brain" (Frege 1884/1960, xviii). As Frege saw things, not only would allowing psychological concerns in foundationalist efforts be counter-productive, it would undermine the objectivity and certainty of mathematics. After all, while ideas about what words like 'tree' and 'number' mean may vary from person to person, the meanings themselves are supposed to be objective: regardless of what we think, 2 plus 2 equals 4. It is no surprise, then, that Frege's review of Mill's mathematical empiricism was particularly scathing, including a passage where Frege describes Mill's views as an 'apocalypse' (Frege 1884/1960, 9). For Frege, numbers are as mind-independent as the water in the North Sea is. If this is the case, then how can psychology teach us anything about them?

In short, the success of Frege and Husserl's anti-psychologistic crusade at the turn of the 20th century did not bode well for naturalistically-inclined philosophers of mathematics – including those interested in Mill and the empirical study of number concepts. In a sense, this dismissal of empirical and psychological matters has been

---

[4] See Kusch 1995, chapter 3, for a detailed presentation.

vindicated: despite the foundational crisis of the early 1900s, mathematics has grown by leaps and bounds and its effectiveness in describing and modelling reality has reached unimaginable heights. Most mathematicians don't care about the foundations of their discipline, since most don't ever encounter problems of the sort identified by Russell and Gödel, for example. Those few who do bother to think about foundational issues are generally platonists on weekdays and formalists on Sundays, as Davis & Hersh (1981) famously put it. From a mathematician's point of view, then, foundational issues that concern the ontological status of mathematical objects is of no consequence, as elegantly illustrated by Morris Kline:

> The developments in this [twentieth] century bearing on the foundations of mathematics are best summarized in a story. On the banks of the Rhine, a beautiful castle has been standing for centuries. In the cellar of the castle, an intricate network of webbing had been constructed by industrious spiders who lived there. One day a strong wind sprang up and destroyed the web. Frantically the spiders worked to repair the damage. They thought it was their webbing that was holding up the castle. (Kline 1980, 277)

There is then a sense in which we can accept that foundational issues haven't been resolved, since this lacuna has had little effect on actual mathematical *practice*.

And yet there is also a sense in which, from a philosophical point of view, the situation is unacceptable. For while it could be argued that concepts like SET,[5] SUCCESSOR and NUMBER have given mathematics solid conceptual footing, the ontological status of these foundational notions remains completely mysterious, as does our epistemic access to these. To provide a true foundations of mathematics, then, foundations of mathematics must widen its area of enquiry to include the practices of real-life mathematicians (cf. Wang 1974, 242-243). This is especially true if we wish to think of mathematics from a naturalist perspective, and thus conceive of it as a (uniquely)

---

[5] Names of concepts and their content are expressed in capital letters throughout the text.

human enterprise with roots in evolution.

So while their campaign was largely successful in eliminating psychologistic tendencies from foundations of mathematics, there are reasons to doubt that the alternatives proposed by Frege and Husserl represent an improvement. Not only did few aspects of their anti-psychologism survive unscathed,[6] their attempts to free mathematics and logic from the limitations of the human mind have been accused of closet psychologism.[7] But more importantly, from a naturalistic point of view, what is especially problematic is that both Frege and Husserl rely on mysterious, platonic entities to replace the more familiar psychological notions they reject.

Of course, a well-known problem with banishing numbers to an abstract platonic realm is that we have no idea how subjective, temporal beings like humans can grasp objective, ideal propositions, since these are independent from time, space, and mind (e.g. Benacerraf 1973). By its very nature, the anti-psychologistic picture seems incomplete with respect to the epistemological question of how we could interact with the objects of mathematical practices. Worse, there is something essentially mysterious, if not esoteric, in postulating the existence of a separate, abstract world outside of space and time.[8] In short, by replacing psychologism with platonism, it looks like Frege was providing "an explanation of the obscure by the more obscure" (Philipse

---

[6] See Kusch 1995 for a review of contemporary (ch.4) and modern (appendix II) criticism of their anti-psychologism. See also Aach 1990.

[7] For example, while Frege talks of primitive and self-evident truths (e.g. Frege 1884/1960, 4) and objective meanings, Husserl speaks of ideal thought content and atemporal propositions. Frege's reliance on self-evident truths seems to involve the difficult notion of a self that is neither subjective nor psychological. Similarly, Husserl's ideal laws are 'intuited a priori'. But what are such intuitions if not psychological processes? Kusch 1995 lists authors who argued Husserl was a closet psychologicist (see also Baker & Hacker 1989, 81-2; Tieszen 1990).

[8] Rav (2007, 87) expresses his views on Platonism thus: "'How is it that the Platonistic conception of mathematical objects can be so convincing, so fruitful and yet so clearly false?' writes Paul Ernest in a review...I disagree with Ernest on only one point: I do not think that Platonism is fruitful. As a matter of fact, Platonism has negative effects on research by blocking a dynamical and dialectic outlook."

1989, 67).

Considering the mysterious nature of platonic entities, it may be wise to give mathematical psychologism another chance. After all, the witch-hunt for psychologism took place at a time when psychology was still in its infancy. As Sober put it: "It is no wonder that Frege conceived of psychological phenomena as variable and erratic, given the way psychology was done at the time" (Sober 1978, 169).[9] Nowadays, given the advances made in various branches of cognitive sciences, the claim that psychological facts are necessarily subjective and mind-dependent seems false. Facts about how brains and minds work seem more and more akin to facts about, say, the water of the North Sea. As neurologist Stanislas Dehaene put it, "With such bridging laws from neurons to behavior, psychology comes increasingly close to being an exact science" (Dehaene 2011, 250).[10]

It could be argued that some attempts at providing purely logical foundations to arithmetic include some form of implicit reliance on properties of real-life mathematicians. Dedekind, who was criticised by Russell and Dummett for importing non-mathematical notions in the foundationalist enterprise is one possible example, while widespread reliance on the axiom of choice could similarly be interpreted along subjectivist lines. But even in these cases, we have no idea how foundational notions like numbers and sets fit into a world in which atoms, brains, hamburgers, and colors

[9] Frege himself seemed to be aware of the embryonic state of psychology at the time, writing that "It would be strange if the most exact of all the sciences had to seek support from psychology, which is still feeling its way none too surely." (Frege 1884/1960, 38). Schlick also seemed aware of these growing pains, but argued that this should not affect the principles behind psychologism: "all processes in nature and mind occur according to laws, and these laws are without exceptions, just like the rules of formal logic. The laws are not inexact, our knowledge of them is insufficient—this is a huge difference" (Schlick 1918, 128).

[10] As Luc Faucher pointed out to me, it is worth mentioning that Dehaene's claim here should be taken with a grain of salt: while it is undeniable that psychology has grown by leaps and bounds, it is still far away from being anywhere as exact as physics, say. Perhaps more importantly, we are also far from having general bridge laws that fit Dehaene's description here.

also figure. By focusing on securing solid concepts on which to base mathematical practice, classical foundationalism has voluntarily kept as far away as possible from such empirical matters.

Thus, while it is true that mathematical *practice* does not need to address empirical and foundational issues in order to proceed, *philosophy* of mathematics certainly does. This is especially true for naturalistically inclined philosophers, *pace* Quine. Indeed, for philosophers, the 'unreasonable effectiveness' of mathematics only adds to the mystery: for not only must we try to understand how abstract and objective mathematical concepts are related to our subjective minds, we must also explain how they are so perfectly suited and effective in describing and predicting events in the material world.

2 New wave foundationalism: the rise of numerical cognition studies

Of course, the fact that mathematical objects have puzzling properties is not new: the ontological status of mathematical entities is one of the oldest questions in the documented history of western philosophy, and platonism towards mathematical entities has long been considered problematic (e.g. see Benacerraf 1965, 1973; Bernays 1935). What *is* new, however, is the emerging body of empirical research into the cognitive foundations of numbers. While research in this field was not widespread until late in the second half of the 20th century (Piaget being a notable, if flawed, exception), things started to change around the time Gelman & Gallistel published their landmark book *The Child's Understanding of Number* (1978). At around the same time, in the sixties in seventies, Lakatos (1976) wrote a series of articles that eventually laid the groundwork for the study of mathematical *practice*, describing the historical development of mathematics as a dialectic process of proofs followed by refutations. Bringing the discussion to the broader implications of mathematical practice, Kitcher

(1984) and, more recently, Ferreiròs (2015), have embraced an interdisciplinary approach to the study of mathematics, framing it as a human activity.[11] So while talking about the relationship between mind and number is still anathema in many philosophical circles, there is hope for the naturalistically-inclined philosopher of mathematics.

This multi-faceted research has led to numerous findings, including the discovery of evolutionarily shared representational systems that could be the source of adult humans' developed arithmetical abilities. When Dehaene (1997) published his classic monograph, *The Number Sense*, interest in the relationship between mind and math was rejuvenated to the point where the cognitive bases of our ability to represent numbers are now the focus of studies in a wide variety of domains of research. Many competing accounts of the relationship between brain and number have tried to make sense of all this data, sparking a new and fascinating debate surrounding the empirical status and developmental bases of numbers. After thousands of years of being in the dark, it looks like science is beginning to shed light on the relationship between world and mathematics.

3 The gap problem

Unfortunately, with the opening of these experimental floodgates comes new questions about the concept NUMBER. The problem, in a nutshell, is this: how do we combine data from developmental psychology, ethnology, linguistics, evolutionary biology, cognitive psychology, neuropsychology, and anthropology to form a mathematically-viable number concept? For example, while the evidence supporting the existence of innate neuronal systems capable of producing imprecise representations of the quantity

[11] See also Van Kerkhove et al. 2006, 2010; Maddy 1996, 1997, 1998.

of objects to which we are paying attention is well established, how this and other systems housed inside brains can be responsible for the objectivity, generality, and efficiency of concepts like NUMBER is still very much an open question.

Putting aside questions concerning whether such results could even meet Benacerraf's (1973) challenge and take us to the promised land of naturalized numbers, a less metaphysically-motivated difficulty we face in answering how such findings relate to abstract objects like numbers is that, in the state they arrive following their long evolutionary journey into our inherited neurological makeup, the psychological systems that figure in most explanations of the material implementation of representations of numbers, including the Approximate Number System (hereafter, ANS) and the Object-File System (hereafter, OFS), are in no shape to produce content with the scope and precision of natural numbers. As I explore in detail in the first two chapters of this thesis, the ANS produces representations that grow increasingly fuzzy as the numerosity it responds to increases (Dehaene 1997/2011), while the OFS likely does not produce any explicit numerical content, and if it did, it would be limited to tracking about four objects (Carey 2009; Kahneman et al.1992). In short, with the discovery of these systems comes the responsibility of explaining what happens to them in order for representations of precise quantities of discrete objects of the sort used in the practice of arithmetic to emerge.

Thus, despite the tremendous progress made by research into the origins of numerical cognition since the late 20th century, an important question remains unanswered: given the limitations of our innate cognitive machinery, how do we manage to build representations with mathematically-viable numerical content? That is, how do we bridge this gap between the quantity-related content produced by our evolutionarily ancient brains and the mathematically-viable numerical content associated with numeration systems like Indo-Arabic or Roman numerals? This is what I will refer to as the *gap problem*. It is the main problem that will concern us in this thesis.

While many have addressed this issue,[12] most proposed solutions rely on the presence of numerical representations in the environment, often in linguistic format, to bridge the gap between the output of our innate representational systems and the content of our advanced number concepts. In this thesis, I argue that this externalist approach needs to be restricted to allow room for one that focuses more on *internal* cognitive processes, since any appeal to external symbols for numbers must come *after* we have explained the emergence of numerical cognition *internally*, given that external symbols for numbers depend on the construction of internal representations with numerical content for their existence.

The main contribution I wish to make with my thesis is methodological: I want to highlight the limitations of externalist approaches to numerical cognition in order to motivate adopting internalist solutions to the gap problem. To do this, I will argue that relying on external representations of numbers in explaining what makes advanced numerical cognition possible leads us to an incomplete – or, worse, circular – account of the origins of our advanced numerical abilities, since externalist accounts have difficulty explaining how numerical content can originate in cultures and environments where there are no external representations for numbers. This is what I will call *the origins problem*. It will orient most of my criticism of externalist approaches. I claim that externalist approaches to numerical cognition cannot explain how we first came up with representations with numerical content – both in our heads and in the world – because they take for granted the very thing we are trying to explain. I will argue that the solution to my origins problem is to restrict the scope of externalism and adopt an internalist approach for an initial segment of the natural numbers.

---

[12] For example, Gelman & Gallistel 1978; Gallistel et al. 2006; Hurford 1987; Dehaene 1997/2011; Butterworth 1999; Lakoff & Nunez 2000; Wiese 2004; De Cruz 2007, 2008; Rips et al. 2008a; Carey 2009; Coolidge & Overmann 2012; Menary 2015a; Malafouris 2010.

4 Thesis outline

In addition to this introduction and the conclusion, my thesis can be usefully divided into three main sections, each containing two chapters. In the first section of the thesis, which includes the first two chapters, I summarize data concerning cognitive systems that are often considered to be the building blocks of numerical cognition – the OFS, explored in detail in chapter 1, and the ANS, in chapter 2 – and describe how their limitations prevent them from providing the full story of the origin of the content our advanced number concepts. This will involve taking a look at data from a wide variety of domains of empirical research as well as their methods. Of course, given the incredible variety of studies that have tackled the empirical foundations of our arithmetical abilities, I cannot list every potentially relevant experiment done in this field. Rather, I will try to limit my treatment to studies suggesting that the ANS and the OFS are present in human adults and preverbal infants as well as in the animal kingdom. If we eventually wish to explain where our numerical abilities come from and how we can represent abstract objects like numbers, it is important to show that the systems that appear recruited in these abilities have evolutionary origins, since this can help explain the origins of numerical content by appealing to Darwinian principles of evolution through natural selection.

The purpose of these initial chapters is twofold: first, in order to set the stage for the gap problem it is important to be clear on what sort of content we can expect to get from our innate cognitive machinery. As will become clear in the first two chapters, there are many complications that arise from the empirical study of numbers, due to their abstract nature. The first two chapters will allow me to set up the gap problem by presenting and evaluating the empirical support for the existence of cognitive systems whose properties make them likely cognitive foundations for our ability to practice arithmetic. By identifying the evolutionarily-inherited cognitive systems that appear recruited in numerical cognition, as well as their limitations, it will become clear that we need an explanation of what happens to these systems in order to explain their use in numerical cognition. The data summarized in the first two chapters will then be used as empirical constraints any account of numerical cognition must satisfy in order to bridge the gap between what these systems allow us to do and what we actually do in the practice of arithmetic.

Such accounts of how we bridge the gap will be presented in the second section of my thesis, which includes chapters 3 and 4. Here, I summarize the main lines of two of the most influential accounts of how we can use the data presented in chapters 1 and 2 to explain the emergence of numerical content. These are Stanislas Dehaene's (1997/2011) number sense, presented in chapter 3, and Susan Carey's (2009) Quinian [sic] Bootstrapping account, which I develop in chapter 4. The main objective of these chapters is to illustrate what sort of strategies have been proposed to explain the precise role of the systems presented in the first two chapters in representing natural numbers, and what sort of transformation on these can explain the development of numerical cognition. These chapters will also highlight the externalist commitments of these accounts of the origins of numerical cognition by identifying the precise role played by objects and symbols outside our heads in the development and practice of numerical cognition.

There are, of course, other proposals on offer. I chose these specific accounts for two main reasons. First, both of these are easily among the most influential in the literature, each author almost invariably popping up in the references section of works dealing with numerical cognition. Second, both these authors offer incredibly detailed and varied empirical support for their accounts, covering data from a wide range of domains. Given the broad range of disciplines that can inform the nature of our arithmetical abilities, this makes them ideal candidates for a philosophical treatment of answers to the gap problem, which requires covering data from many disciplines. While Dehaene's work could arguably be described as focusing more on the neurological aspects of our arithmetical abilities, Carey's can be described as focusing on the developmental side of this topic, including critical data obtained by Karen Wynn (1990, 1992) on the stages children go through when learning the meaning of number words. In this sense, although both accounts cover data from a wide range of domains, they are complementary approaches.

These models are also opposites, in a sense, given that Dehaene attributes a central role to the ANS in his account, while Carey argues that the ANS only comes into the picture once the OFS has done the hard work, along with other systems that I discuss in chapter 4. As will become clear, despite the differences in Dehaene and Carey's approaches, both crucially rely on the presence of external symbols and artefacts with numerical content to explain both the historical and developmental origins of our ability to practice arithmetic. The fact that both can be considered externalists about the development of numerical cognition despite considerable differences in their approaches illustrates the variety of externalist answers to the gap problem and their shared commitments. At the end of chapter 4, I discuss my main issue with externalism (i.e. the origins problem described above) and show that it applies to both accounts. This sets up the last section of the thesis, which comprises chapters 5 and 6.

In these last two chapters, I explore potential externalist replies to my origins problem. These rely on considerations related to attributing a constitutive role to external artefacts in numerical cognition, explored in chapter 5, and explaining the development of numerical content by appealing to cultural evolution and the constructive power of our cultural environment, discussed in chapter 6. The first reply to my origins problem involves setting up the motivation for adopting an externalist framework. I do this in chapter 5 by sketching the main lines of Andy Clark and David Chalmer's classic 1998 article, 'The Extended Mind', easily considered the locus classicus of active externalism. This is followed by a brief summary of Clark's (1998) explanatory criteria for extended cognitive systems, which I will use to evaluate to which extent adopting an externalist approach to numerical cognition can bridge the gap. I then illustrate how this externalist framework has been applied to numerical cognition by summarizing Catarina Dutilh Novaes' (2013) analysis of three levels of constitutivity of external representations of numbers and the associated claim that numerical content does not emerge in our head. In my criticism of this account, I take a look at the fascinating case of anumerate cultures like the Pirahã (Frank et al. 2008) and the Mundurucu (Pica et al. 2004) as well as data concerning the development of external representations for numbers in Sumeria around 5000 years ago (Malafouris 2010; Schmandt-Besserat 2010), in order to characterize the conditions in which numerical content can emerge. My main contention in this chapter is that externalist accounts fail to explain what makes the difference between an individual that has bridged the gap and one who has not in situations where both have access to the same external support for cognition, since the difference in these cases lies squarely in our head.

Last, in chapter 6, I consider replies to this internalist criticism by exploring the potential contribution of culture to the gap problem. In the first section of this chapter, I take a look at Helen De Cruz' (2007) Darwinian approach to numerical cognition, which applies elements of cultural evolution to the historical development of mathematics, in order to determine to which extent appealing to mechanisms of cultural evolution can help the externalist bridge the gap. I argue that the problem with appealing to mechanisms of cultural evolution to explain the development of novel content is that these mechanisms operate at the level of populations, while the novel content emerges inside individuals' heads.

This is followed in the second section of chapter 6 by a presentation of Richard Menary's (2015a) enculturated approach to numerical cognition, according to which cognition extends into our cultural niche. Here, I argue that while Menary's enculturated framework can help explain what makes the difference between numerate and anumerate *cultures*, it cannot help specify what makes the difference between numerate and anumerate *individuals*. I argue that enculturation does not have an account of innovation capable of explaining how individuals manage to improve and modify the practices of their cultural niche. This is because enculturation focuses mostly on the inheritance and transmission of practices, not on their origins, which involve individual-level understanding, rather than population-level practices and social pressures. The upshot is that culture provides the necessary background conditions against which individuals can innovate. This role is crucial in the development of numerical abilities – crucial, but explanatorily limited.

5 Terminological digression

Before starting with chapter 1, it may be useful to clear up potential terminological issues that could arise given the multidisciplinary nature of the topic covered in this thesis. After all, the abstract nature of numbers means that one must tread particularly carefully when trying to determine their relationship with material things like human brains. This is because numbers are a unique object of study, in that they appear to be as mind-independent as it gets, and yet, unlike rocks and other mind-independent things, they also appear to be a human creation. De Cruz et al. (2010) describe this peculiar aspect of numbers well:

> numbers are not just abstract entities that are subject to mathematical ruminations—they are represented, used, embodied, and manipulated in order to achieve many different goals, e.g., to count or denote the size of a collection of objects, to trade goods, to balance bank accounts, or to play the lottery. Consequently, numbers are both abstract and intimately connected to language and to our interactions with the world. (De Cruz et al. 2010, 59)

Given that the philosophical study of numerical cognition lies at the intersection of so many disciplines, it is important to attempt to clarify, as much as possible, what it is we are talking about, and which aspects of numbers and cognition we wish to describe.

For while it may be acceptable in certain disciplines to talk of animals having number concepts and preverbal infants doing arithmetic, such loose usage of central terms can lead to confusion when trying to pinpoint exactly what sort of properties we can attribute to human and animal behavior that seems to result from variations of discrete quantities of objects in the environment.[13] While efforts will certainly be made in the following text to introduce any technical terms as they come up, a few central one are worth discussing at the outset.

First, given that a central topic here is the origins of numerical cognition, it is important to qualify what is meant by this expression. After all, as Nunez (2017) has argued, the literature on numerical cognition is polluted by loose and inaccurate use of important terms, especially all things 'numerical'. In an attempt to limit my contribution to this unfortunate trend, throughout this text, I will strive to use the general expression 'numerical cognition' to describe only the type of cognition observed in the practice of formal arithmetic, as done by people who have mastered numeration systems like the Indo-Arabic numerals and the physical manipulation of abacuses. Numeration systems will be considered systems containing rules for the manipulation of external representations for numbers. Here, the term 'symbol' will be used as a synonym for representation. So, on this reading, there can be symbols and representations for numbers in the world, for example, in the form of written numerals or spoken number words, and there can be internal representations and symbols for numbers, in the form of patterns of brain activity. While I will allow myself the luxury of using the term 'numeral' to describe any external representation of number for stylistic simplification, there will be cases where it will be relevant to distinguish these from number words, tallies, and other external symbols for numbers. In such cases, the term 'numeral' will be restricted to written symbols for numbers, as implemented by Indo-Arabic or Roman

---

[13] See Núñez 2017.

numerals. The context should make it clear whether or not the term is being used in this restricted sense.

The difficulty of knowing where to draw the line between those that understand what numbers are and those that do not is well known (e.g. Rips et al. 2008a; see also Relaford-Doyle & Núñez 2018, Núñez 2017). I will often talk about levels of *arithmetical* abilities to describe levels of mastery of formal rules of arithmetic. ‘Numerical abilities’ is a more general expression here, and can refer both to the ability to respond to variations of numbers of objects in our environment, which is shared by infants and animals, as well as to a variety of levels of mastery of numeration systems. On the other hand, ‘numeration skills’ will be used to describe the human ability to enumerate collections of objects. Thus, numeration skills are a type of numerical ability here.

Strictly speaking, then, numerical cognition is only something educated adults can do. However, to reflect different stages of mastery of numeration systems – or lack thereof – I will often qualify the expression ‘numerical’ to emphasize to which extent an individual can be described as being proficient in the practice of formal arithmetic. For example, ‘formal numerical cognition’ and ‘formal numerical practices’ will be used to highlight mastery of the practice of formal arithmetic using numeration systems like the Indo-Arabic numerals, while ‘rudimentary’ numerical cognition will describe individuals who are in possession of the building blocks of formal numerical cognition, including the ANS and the OFS, but who are not capable of distinguishing between collections of arbitrary sizes based on the number of items they contain. In this sense, animals and preverbal infants can be considered as having the rudiments of numerical cognition, or as being able to perform rudimentary numerical cognition, since they possess cognitive systems that react to variations of numerical information in their environment. but they have not developed the ability to distinguish between collections[14] of arbitrary sizes based on the number of items they contain. However, they do *not* possess what I will often call 'proto-numerical abilities' nor formal numerical abilities.

Rather, I will reserve the term 'proto-numerical' to refer to what could be considered an intermediate stage in the development of numerical cognition. At this stage of development of numerical cognition, individuals have the ability to accurately count beyond the subitizing range (i.e. beyond four objects), but their ability to label the numerosity of collections does not include mastery of the generative syntax of a formal numeration system capable of labelling the numerosity of reasonably large collections (see Schlimm 2018 for why it would be inaccurate to speak of 'arbitrarily' large collections here). For example, as I detail in chapter 5, neither the Mundurucu (Pica et al. 2008) nor the Pirahã (Gordon 2004; Frank et al. 2008) would be considered as having proto-numerical skills here, since it is debatable whether the Pirahã have labels for numbers at all (Frank et al. 2008) while the Mundurucu lexicon only contains words for the first five numbers, and these are not used consistently (Pica et al. 2008). Individuals from cultures like the Oksapmin, whose body-part-based numeration system stops at 27 (Saxe 1981), could be considered as having proto-numerical abilities, as do children who have mastered the ability to use number words to accurately label the number of objects in collections containing more than four objects. Individuals from cultures like ours that master numeration systems that can be extended indefinitely will be described as having advanced, fully-fledged, or developed numeration skills or abilities.[15]

---

[14] To avoid confusion with the technical notion of a set as used in set theory, I will try to avoid speaking of groups of material objects using this term, and will instead follow De Cruz et. al (2010) as well as Schlimm (2018) in adopting the more neutral term 'collection'.

[15] Both advanced and proto-numerical skills could further be distinguished from the non-numerical skills of individuals and cultures like the Pirahã (Gordon 2004; Frank et al. 2008) that seem to entirely lack any form of explicit representation of discrete quantities. For those cultures like the Mundurucu (Izard

While there is good reason to think that there have been many degrees of complexity and expressive power of systems of representations of discrete quantities throughout history (e.g. Ifrah 1998; Chrisomalis 2010; Menninger 1969), it would appear pointless to distinguish numeration systems based on the exact number of items they have developed symbols for. For example, it would seem useless to distinguish between systems that have symbols for up to 27 items, and systems that have symbols for up to 29 items, say. To describe the expressive power of such relatively limited systems, I will use the same term as I use to describe the skills of those individuals that master them, i.e., 'proto-numerical'.

As for motivation behind drawing the lines as I have here, the idea here is to reflect different stages of understanding of what numbers are. These distinctions need not be considered as being set in stone but can hopefully facilitate the discussion. As is usually the case, the context in which terms are used should clarify in which sense they are being used.

One arguably important aspect of numbers that will come up here is that they are representations of quantities of *discrete objects*, as opposed to quantities of continuous magnitudes like size or luminosity, as will become clear in chapter 2. While it may be argued that understanding that numbers can be extended indefinitely is a key aspect of numerical abilities, I tend to favor the opposite view, according to which what matters in attributing numerical abilities is the extent to which an individual can distinguish collections based on the number of discrete objects they contain, not their understanding of the infinite nature of the natural numbers. The main explanatory challenge I want to tackle here is that of explaining the ability to distinguish collections containing sufficiently large numbers of objects based on the precise number of discrete

---

et al. 2008; Pica & Lecomte 2008)) who have words for discrete quantities, but not beyond the subitizing range (up to four objects), we could reserve the term pre-numerical, though this level of distinction will not be required here.

objects they contain, given that no innate cognitive system allows this. In this sense, once we have bridged the gap, the ability to extend our numeration abilities seems like a separate, though clearly related matter. A similar thought could be behind Crossley's comment that "once the idea of counting has emerged, then the idea to go on counting does not seem to lie far below the surface" (Crossley 1987, 13).

As I see it, the transition from rudimentary to proto-numerical cognition is the most important developmental step, since it is this step that allows individuals to determine the precise number of items in collections beyond the subitizing range, even if this ability does not extend indefinitely, as it might for people who have mastered formal numeration systems like the Indo-Arabic numerals. The bottom line for me here is that if individuals can count, calculate, and trade using their symbols for discrete quantities, they should be considered as having some form of understanding of what numbers are – or at least, more so than individuals who cannot count or calculate at all – even if they can't easily tell the difference between 1000 and 1001, for example. The gap problem I want to talk about in this thesis is how individuals get from this absence of an ability to precisely quantify to such proto-numerical skills. For this purpose, knowing how we manage to extend our understanding of what numbers are into the realm of the infinite seems like a separate matter from how we manage to build representations of precise quantities in the first place.[17] To understand what sort of tools we are equipped with to deal with this problem, I now turn to the literature on representational systems that appear to underlie our ability to represent quantities of discrete objects.

---

[17] Indeed, there is evidence that it takes roughly two years for children to learn that all numbers have a successor once they can use number words accurately (Cheung et al. 2017), while understanding the infinite divisibility of the number line seems related to mastery of notation for fractions (Smith et al. 2005).

CHAPTER I

THE OBJECT-FILE SYSTEM

1 Introduction

For the naturalistically-oriented philosopher, this is a great time to be interested in the foundations of mathematics. Contrary to the foundationalist approaches of the early twentieth century, whose focus was exclusively on providing formal tools capable of serving as secure, contradiction-free conceptual foundations for mathematics, it is now acceptable to inquire into the epistemological question of how material beings like humans can come to know anything about abstract objects like numbers. In direct opposition to the rampant anti-psychologism of traditional approaches to foundations of mathematics, numerical cognition is now a thriving field of scientific enquiry. By looking for behavioral and activation patterns associated with tasks that require processing information about discrete objects in our environment – i.e. numerical tasks – in both symbolic and non-symbolic format, researchers from a wide variety of fields are starting to uncover the distinguishing properties of the representational systems underlying behavior in such tasks, thereby shining light on the possible origins of our advanced numerical abilities. This is good news for the naturalist, since it finally opens the door for a concentrated, well-supported alternative account of how numbers could be the product of human minds that does not share the mysticism and disregard for actual psychological processes that plagued intuitionism. In short, there has never been a better time to be interested in the way our brain constructs our advanced numerical abilities.

Of course, it would be folly to claim that there is a straight road from numerical cognition data to psychologism in mathematics. For while this road is certainly a tempting one for many – especially, perhaps, those coming from the scientific side of the debate – there is still considerable philosophical work to do before cognitive scientists can even start a discussion with philosophers of mathematics to offer them a viable alternative to platonism and the anti-psychologism that comes with it. In fact, some philosophers have even gone the opposite route, using numerical cognition data to prove that platonism in mathematics is alive and well (De Cruz 2016).

In this chapter, I take the first steps on the road towards providing a naturalistic account of the psychological origins of mathematically viable number concepts by summarizing findings concerning one of the two main candidate cognitive systems underlying our basic numerical abilities, so that I can then sketch how these systems could produce representations with the abstract character of numbers in later chapters. In section 1.2, I survey findings supporting the existence of two distinct systems underlying our ability to enumerate, including serial counting and subitizing, and I argue that the best way to explain subitizing is via Trick and Pylyshyn's (1994) visual indexing model. Then, in section 1.3 I describe the relationship between these visual indexes and object files, summarizing evidence for the existence of an object file system in infants and animals. I close the chapter by discussing to which extent the evidence supports attributing numerical content to this system.

1.2 Evidence for two Systems

1.2.1 Subitizing

While numerical cognition has only grown into a coherent and widespread area of research in the past thirty years or so, at least one aspect of our numerical abilities has been studied for more than one hundred years. In the late 19th century, Jevons (1871) and Cattell (1886) observed that adult humans' ability to accurately label the number of objects perceptually presented in a visual array – its numerosity[18] – dropped sharply when this number rose above 3. Starting at 4 stimuli, enumeration errors started creeping in, confidence in answers dropped and reaction time increased linearly in relation to the number of objects displayed thereafter. This ability to quickly and accurately apprehend the numerosity of small collections of less than four objects was eventually dubbed *subitizing* (Kaufman et al. 1949).

In typical distractor-free subitizing experimental paradigms, each stimulus within the subitizing range (1-4 stimuli) increases response time by 40-100 milliseconds, while outside of this range, each adds an extra 250-350 milliseconds to the response time (Trick & Pylyshyn 1994, 80). Enumeration is fast, effortless, and accurate, when the number of objects is kept below 4.[19] Confidence in numerosity judgements in the

---

[18] The term 'numerosity' has unfortunately come to mean different things to different people in recent years, some authors (Malafouris 2010; Coolidge & Overmann 2012) having taken the surprising decision to use it to describe the cognitive system that detects quantities of discrete objects in collections instead of these collections themselves, in opposition to what was common practice in the numerical cognition literature (see Schlimm 2018, 197). In what follows, I will use the term 'numerosity' to describe the number of objects an observer perceives a collection as having. This perceiver-relativity of numerosity emphasizes that numerosities are not numbers, since they are physically-detectable collections of discrete objects, not abstract entities of arbitrary size.

[19] Some have reported individuals whose range goes as high as 7. See Trick and Pylyshyn (1994) for why such individual differences need not matter for identification of the underlying cognitive systems. Methods requiring participants to enumerate via pointing may increase the subitizing range to 6 (Haladjian & Pylyshyn 2011). There is also evidence that it is possible to increase one's subitizing range, e.g. by playing lots of video games (Green & Bavelier 2003, 2006), though there is reason to doubt to

subitizing range is also higher than that of the counting range. On the other hand, when enumerating objects outside this range, the process is slower, error-prone, and requires conscious effort, either in the form of counting the objects serially one-by-one, or grouping them together into small sets and then adding the total number of clusters to form an approximate estimation of their number. Even when experimental manipulations of stimuli slow down enumeration, they generally affect processing both inside and outside the subitizing range, so that the effect remains visible, as evidenced by the tell-tale elbow-shaped slope discontinuity in graphs plotting response times versus number of items in subitizing studies (see Figure 1).[20]

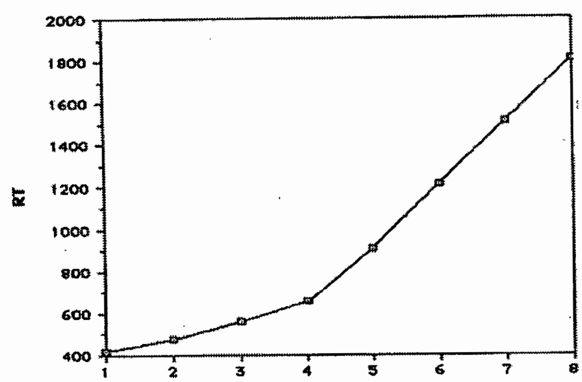

1

Figure 1.1 The subitizing elbow. (From Trick & Pylyshyn 1994)

---

which extent the evidence is strong here (Guha et al. 2014).

[20] Reaction times actually may not follow such straight slopes in the subitizing range, since the reaction time difference between 1 and 2, 2 and 3, and 3 and 4, can all differ (Dehaene & Cohen, 1994). This led some to deny the existence of a distinct status for enumeration in the subitizing range (Balakrishnan & Ashby 1991).

The evidence for this behavioral discontinuity in enumeration tasks is robust. The question is, how do we explain it? Following Jevon's and Cattell's groundbreaking work, researchers have been trying to figure out how subitizing works, why it is limited to such a small numerical range, and whether or not it recruits the same cognitive system involved in enumerating larger quantities of objects. While many questions remain, considerable progress has been made, as I hope to show in the next few sections. First, I take a look at evidence supporting the existence of two distinct systems for subitizing and enumeration of larger quantities of objects (section 1.2.2). Then, I discuss some of the most promising explanations of how subitizing works, in order to identify the cognitive system that underlies this ability (1.2.3, 1.2.4). I then take a closer look at how this system relates to other systems involved in the process of enumeration in the subitizing range(1.3), to find out what sort of numerical content, if any, is associated with it (1.4).

1.2.2 Two paths to enumeration

There has been considerable debate over the years regarding whether the behavioral discontinuity in enumeration processes is mirrored at the cognitive level, or if a single system underlies behavior inside and outside the subitizing range.[21] The issue here is whether the cognitive systems underlying subitizing are distinct from those that allow estimation of larger collections of objects. On the single-system side, while some (e.g. Gallistel & Gelman 1991, 1992, 2000; Cordes et al. 2001) proposed that subitizing is a form of rapid counting, others (e.g. Dehaene & Changeux 1993; Dehaene 1997/2011; van Oeffelen & Vos 1982; Vetter et al. 2008) argued that it is a form of rapid estimation.

---

[21] For a short, fair summary of both sides, see Chesney & Haladjian 2011. See also Trick & Pylyshyn 1993, Gallistel 1990, Balakrishnan & Ashby 1991, Dehaene 1997/2011, Piazza et al. 2003, 2011, Feigenson et al. 2004, Ansari et al. 2007, Deymeyre et al. 2010, 2012.

There is behavioral and brain evidence to support this claim. For example, imaging data obtained from Positron Emission Tomography (PET hereafter) readings of adults were taken as they identified the number of dots on a screen both within (1-4) and above (6-9) the subitizing range, and found that a common network was activated in both ranges, although this network is more active in the counting range. They interpreted this as indicating that "Subitizing does not seem to rely on a separate dedicated neural mechanism that is not also involved in counting" (Piazza et al. 2002, 442).[22]

However, while the question of whether a single cognitive system can explain behavior both in the subitizing range and for larger numerosities is still open, single-system theories have now fallen out of favor (Dehaene 1997/2011; Feigenson et al. 2004; Vetter et al. 2008) in the face of mounting behavioral and neuroimaging evidence of differences in behavioral signatures between the subitizing and counting range. Nowadays, a majority of authors adopt a two-systems approach to numerical cognition.

In a testament to just how fast numerical cognition studies are developing, a number of authors who once advocated the single-system approach have switched sides. For example, in the second edition of his 1997 classic, *The number sense*, Stanislas Dehaene explicitly recants the one-system position adopted in the first edition, following a study in which reaction times for quantity estimation tasks were compared in relation to the ratio of the quantities displayed (Revkin et al. 2008).[23] The idea here

---

[22] See also Cordes et al. (2001) who found similar behavioral signatures (scalar variability) inside and outside the subitizing range when participants were prevented from counting verbally while tapping a key a specific number of times. Arp et al. (2006) conducted experiments with children affected by cerebral palsy and found that subitizing performance correlates with enumeration performance. See Demeyre et al. 2010 for a balanced discussion of support from lesion studies.

[23] Similarly, while Piazza and colleagues found evidence of a shared network for subitizing and counting, two of the authors involved in this publication (Manuella Piazza and Brian Butterworth) later published evidence supporting a two-systems approach (Piazza et al. 2003, 2011; Agrillo et al. 2012).

is that if only one system were responsible for our performance in both the subitizing and the counting range, we would expect behavior to follow the same performance signature, regardless of how many objects needed to be enumerated. For example, if a single system were involved, the difference in time between identifying four objects and one object should be proportional to the difference in the time it takes to identify 40 objects compared to 10 objects, given the same ratio in both cases. This does not appear to be the case.

Rather, the speed with which we identify numbers of objects within the subitizing range is consistently faster than any other range, even when matching for ratios between quantities of objects. Adults asked to identify the number of dots presented on a screen are markedly faster when the number is kept below 4-5, even when taking into consideration these ratios. According to Dehaene, these results "leave no doubt that a distinct process deals with the subitizing range of numbers—a conclusion that has also been supported by brain imaging research" (Dehaene 2011, 257).

Evidence from neuroimaging studies does indeed appear to favor the two-systems approach, even though some data is consistent with single-systems interpretations, as mentioned above. Consider the neuronal resources required for transitive counting (i.e. counting objects, as opposed to reciting a numeral list). This task requires orienting attention to each object in order. Such serial orienting of attention recruits posterior parietal networks, so if these networks are recruited in counting, but not in subitizing, this counts as evidence that the two processes are separate. To test this, Piazza et al. (2003) took functional Magnetic Resonance Imaging (fMRI hereafter) readings during both quantity identification and color naming tasks to determine whether separate networks were activated by subitizing and counting. They found that posterior parietal activation was much more active during counting when compared to subitizing, suggesting a separation of parallel preattentive processing in subitizing, and a serial attentional processing for counting.

As a sign that neuroimaging data on this are still inconsistent and in need of further confirmation, Ansari and colleagues (2007) found opposite dissociation patterns: they found increased activation in the right temporo-parietal junction for magnitude comparisons of non-symbolic stimuli in the subitizing range, while activation of this area was suppressed outside the subitizing range. This area is associated with stimulus-driven attention (Corbetta & Shulman 2002), suggesting that attention is involved in processing of small numerosities in the subitizing range, but not in the counting range.[24] Piazza and colleagues later found evidence that individual differences in subitizing reflected individual differences in working memory, while this did not hold true for estimation abilities, further supporting the two-systems approach, and indicating that subitizing may recruit a domain-general mechanism (Piazza et al. 2011; see also Ashkenazi et al. 2013).

Further, Hyde & Spelke (2011) took Event-Related Potential (ERP hereafter) readings from the scalp of young infants (aged 6-7 months) while they were exposed to displays on which groups of objects within and beyond the subitizing range were presented. Replicating results obtained in adults using the same method (Hyde & Spelke 2009), they found evidence that processing of numerical stimuli is separate for the subitizing range from a very young age, suggesting that this dissociation is innate, rather than being the result of culture or learning. There is also evidence of dissociations in subitizing and counting in subjects suffering from brain lesions (Demeyere et al. 2010, 2012; see also Cipolotti et al. 1991). In short, the list of behavioral and brain studies that support the existence of different cognitive mechanisms for subitizing and counting is very long, so while the matter has not been settled definitely, there is at this point considerably more support for the two-systems model.[25]

---

[24] This result was later confirmed and expanded upon by Burr et al. (2010). As I discuss below in section 1.2.6, there is disagreement concerning the role of attention in subitizing.

[25] For example, Watson et al. (2007) show that there is a sharp increase in eye saccade frequency above

Thus, considerable evidence mitigates against one-system explanations of the performance discontinuities in numerosity-estimation tasks. Rather, the speed and accuracy with which quantities are identified[26] and the effortlessness of this process when compared to enumeration outside the subitizing range appear to depend on two distinct cognitive systems. Importantly, this two-systems approach means that numerical cognition recruits a patchwork of cognitive systems rather than relying on a single number-dedicated system.

To get an idea of where numerical abilities come from, then, we must consider what sort of cognitive systems are behind both subitizing and quantifying objects outside the subitizing range, and how these systems allow the development of advanced numerical cognition. This will involve finding out why the subitizing range is so restricted and why we treat small and large numbers differently: what is special about the number four, and why are there behavioral differences between enumerating small and large numbers of objects? In the remainder of this chapter, I sketch the most widespread explanation for subitizing and review evidence supporting its existence in adults, infants, and animals. Then, in chapter 2, I summarize the main characteristics of the system underlying basic numerical abilities outside the subitizing range in humans and animals.

---

the subitizing range and that subitizing performance remains fast and accurate despite restrained eye movement. Researchers have also found that anxiety affects counting but not subitizing (Maloney et al. 2010). For more recent neurological evidence of divorced processing, see Fornaciai & Park 2017, who develop a three-system model to account for reduced enumeration performance when displays contain so many dots they appear cluttered. See also Simon 1997; Trick & Pylyshyn 1994; Ansari et al. 2007; Agrillo et al. 2012; Lipton & Spelke 2003; Cutini et al. 2013; Burr et al. 2010; Pagano et al. 2014; Sathian et al. 1999 for more evidence supporting the two-systems approach.

[26] As Trick and Pylyshyn (1994) show, subitizing is not always fast, since some enumeration tasks require more time to register and group perceptual features. It is however generally faster than enumeration outside the subitizing range.

1.2.3 Subitizing as pattern recognition

A number of explanations have been proposed to explain enumeration abilities in the subitizing range over the last one hundred years.[27] A frequently cited model of the nature of subitizing is Mandler & Shebo's pattern recognition account (1982).[28] To explain the speed and accuracy of subitizing, Mandler & Shebo rely on the fact that one dot makes a point, two dots make a line, and three non-aligned dots draw a triangle. On this model, we would easily identify a dot array as containing three items because of our ability to recognize triangular arrangements of objects, regardless of their configuration. Our ability to quickly and effortlessly identify the number of objects when it is kept low is explained by the fact that we are familiar with the patterns that the arrangements of dots make up, which makes it easy for us to identify the number of objects they comprise.

Mandler & Shebo explain the sudden drop in reaction time after four by noting that there are few canonical patterns of stimulus combinations formed of more than three objects. For example, all triangles have three points, but four points can make up a rectangle, a losange, a square, a trapezoid, etc. One source of support for this view is the fact that Mandler and Shebo (1982) have shown that performance speeds up outside the subitizing range when objects in displays are arranged in familiar patterns, like the sides of a die. Further support for this approach comes from the finding that children suffering from developmental dyscalculia present deficits in both subitizing and counting ranges and do not display increased reaction times when dot arrays are

---

[27] See Trick and Pylyshyn 1994 for a critical review of many proposals.

[28] See also Kaufmann & Nuerk 2008; Ashkenazi et al. 2013; Gliksman et al. 2016, for more recent discussions of pattern recognition and subitizing.

presented in canonical form (Ashkenazi et al., 2013), suggesting that pattern-recognition and subitizing may recruit the same resources. Further support comes from Gliksman et al. (2016), who showed that enhancing attention (e.g. via cuing or alerting) produces subitizing-like behavior outside the subitizing range, by enhancing pattern-recognizing global processing.

Despite this support, an important problem with this view is that patterns do not always co-vary with number. For example, it is possible to arrange any number of objects on a line, which should imply that any display containing objects on a line would be slower to process, or inaccurate, due to the interference of the pattern associated with two stimuli. Similarly, it is possible to arrange multiple objects in quasi-triangular fashion, by placing dots close to angles, which should elicit the response "three". Appeal for this model is also undermined by the fact that some research has found activation of object recognition areas irrespective of pattern of presentation, and no evidence of differences in neuronal networks activated by stimuli presented in canonical form vs random arrangements (Piazza et al. 2002).[29]

1.2.4 Trick and Pylyshyn's visual indexes

Taking their cue from Zenon Pylyshyn's (1989, 2001) work on visual indexing, Trick and Pylyshyn (1994) have proposed another well-received model according to which subitizing occurs as a by-product of the way in which visual objects are individuated. According to this visual indexing theory, subitizing rests on a limited capacity cognitive system that automatically creates a small number of indexes of visual feature clusters based on their spatiotemporal continuity, allowing us to keep track of what

[29] For more on the limitations of this view, see Trick and Pylyshyn (1994).

went where in our visual field. While these indexes themselves do not encode the properties of the objects being tracked, they do pave the way for further processing in line with the demands of the task at hand by selecting some features out of the visual field.

The main purpose of this indexing system is to tag a small number of perceptual features so as to give them higher processing priority for attention. Visual indexing "provides a means of setting attentional priorities when multiple stimuli compete for attention, as indexed objects can be accessed and attended before other objects in the visual field" (Sears & Pylyshyn 2000, 2). It is important to note that this indexing system displays the same capacity limits as subitizing: only a limited number of visual indexes can be individuated simultaneously.[30]

Pylyshyn devised many Multiple Object Tracking (MOT hereafter) paradigms to confirm the existence of such an individuating mechanism and show that its capacity limits correspond to those of the subitizing range.[31] In a typical MOT trial, visual stimuli (e.g. geometrical shapes) are presented on a screen and participants' attention is directed towards a particular group of objects (e.g. by having them flicker on the screen) to mark them as targets that they will have to track. Then, the stimuli move around the screen in random motion and the subjects are tasked with correctly identifying them once they stop moving at the end of the trial. Typically, the target stimuli move among other identical non-targeted objects that act as distractors to the task (see Figure 1.2).

In MOT, changing features of target stimuli (e.g. color, shape, size) or even their kind

---

[30] Yantis proposed a similar system in charge of assigning tags for processing priority, based on findings indicating that the number of stimuli that can capture attention simultaneously is limited to four objects (e.g. Yantis & Johnson 1990).

[31] See for example Pylyshyn 1989, 2003; for a Multiple Object Tracking review see Scholl 2009.

(e.g. a frog changing to a prince) does not limit our ability to track them, which suggests that individual features of objects are not required for tracking. Rather, what is important is spatiotemporal continuity of feature clusters. Evidence for this comes from the fact that if objects change features during motion (e.g. constant change of color), when the motion stops, if an object disappears from the display, subjects can identify its last location, but not the relevant feature (Scholl et al. 1999). This indicates that spatiotemporal continuity is processed at a higher priority level than individual features. Moreover, if objects adopt strange trajectories that take them in and out of existence in ways that do not correspond to passing behind an occlusion (e.g. by shrinking to nothing on one side of the occlusion and reappearing on the other side), tracking fails (Scholl & Pylyshyn 1999). This suggests that the system allowing us to track multiple objects distinguishes occlusion from non-existence.

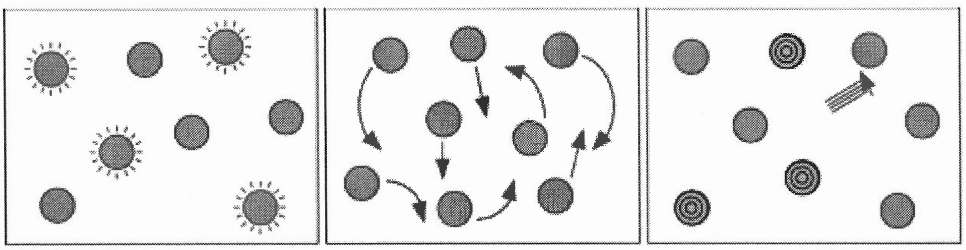

2

Figure 1.2 Multiple-Object Tracking. (From Pylyshyn 2003)

The important point to note about behavior in MOT is that it often displays high success rates for tasks involving less than 5 objects, thus displaying the same capacity limits observed in subitizing. While the precise reason for this capacity limit is not known, it is often associated with limits to working memory (see Pagano et al. 2014) or object-directed attention (e.g. Mandler & Shebo 1982; Scholl 2001).[32]

Such capacity limits in parallel individuation of visual indexes can explain why subitizing is limited to a few objects: since the construction of object indexes occurs in parallel, we can easily identify the number of objects in the subitizing range, given that this does not require any serial attentive processes like those involved in counting, such as marking, indexing, and remembering which objects have already been counted. However, once the number of objects exceeds the subitizing range, attentional resources must be used to count the objects serially, which requires more time and processing. On this model, then, difference in response time between subitizing and counting is attributable to the taxing effect of serial attention in counting.

According to Trick and Pylyshyn, subitizing occurs in two stages. In the first, 'prenumeric' stage, visual indexes are assigned in parallel to each stimulus in the display. It is only in the second stage, dubbed 'number recognition', that numerical content comes in, when subjects choose a response from their numerical lexicon. This second part of the process does not occur in parallel, since words for numbers must be retrieved from memory, where, according to Trick and Pylyshyn, they are stored in order. Following the work of Klahr (1973), they claim that numerical recognition

---

[32] The fact that visual indexes share features of both attentive and preattentive processing can provide some insight into why this limit would be at four objects. Like many preattentive processes, visual indexes are individuated in parallel. Unlike other preattentive processes, in some circumstances, they are sensitive to goal-directedness: subject's intentions can modify the way in which objects are individuated (Trick & Pylyshyn 1994). On the other hand, like attention, the individuation of visual indexes is capacity limited. Such attention-constrained capacity limits of reference tokens pre-attentively is to be expected, since there would be no benefit in selecting more objects pre-attentively than can be handled by attention. And yet, there is also a reason why a small number of objects must be individuated in parallel, since this is the only way for attention to be able to orient itself spatially through the visual landscape. Otherwise, if only one object were individuated, attention could not orient itself towards the next object to be counted. On the other hand, if too many objects are individuated, the processing costs become high. Thus, having a small number of objects individuated in parallel is a good way to balance being able to compute spatial relations between these and orient attention accordingly without creating too large of a cognitive load.

involves "matching each individuated item with a number name, in the order of the number name" (Trick & Pylyshyn 1994, 88). This is what explains the slight slope in response time for the subitizing range: word retrieval takes a little longer for each successive number word, since the matching from index to number word starts at 1.

1.2.5 Support for parallel visual indexing

Support for such an object-individuating explanation of subitizing comes from many experimental sources. One particularly compelling source of behavioral evidence for this explanation involves presenting visual stimuli in the form of retinal afterimages produced by flashguns.[33] Since the perceived dots are not actual physical objects but retinal afterimages, serial visual tagging of these is prevented, as the eyes cannot move from one object to the next to individuate them in a serial counting routine.[34] When stimuli are presented in this manner, there are high error rates in quantity identification tasks outside of the subitizing range, even when afterimages persist for up to a minute, suggesting that failure to individuate objects serially prevents precise enumeration of quantities outside the subitizing range. As predicted, accurate enumeration is preserved for small numerosities, since no such serial processes are required in the subitizing range, given that we can individuate a limited number of objects in parallel and simply read off their number without having to attend to each serially.

Further support comes from a study where participants were asked to simultaneously

---

[33] See Atkinson et al. 1976, replicated and expanded in Simon & Vaishnavi 1996.

[34] In typical subitizing studies, subjects are prevented from counting by restricting the time during which stimuli are exposed. This method however leaves open a single-mechanism explanation of subitizing according to which subitizing is a form of very fast counting that can occur in short temporal windows. Using afterimages prevents such counting to be considered a possible alternative explanation.

track multiple objects and subitize (Chesney & Haladjian 2011). If, as Trick and Pylyshyn proposed, subitizing recruits individuation mechanisms that allow MOT, then interference between these two tasks should be evident. As predicted, results showed that the number of objects in participants' subitizing range decreased in direct relation to the number of objects they were asked to track, confirming that parallel individuation of visual objects is a limited resource shared by both MOT and subitizing.

Trick and Pylyshyn also confirmed that subitizing depends on individuation processes via a series of behavioral studies that tested the effects of feature arrangements on subitizing performance. The idea here is that if subitizing depends on automatic individuation, when such individuation is blocked, subitizing shouldn't occur. One way to prevent such individuation from happening is to confound object boundaries, making it difficult to create visual indexes for these. Thus, paradigms in which the stimuli are arranged in a way that leads to grouping together features from different objects should affect subitizing, since the underlying feature grouping that accounts for the rapid response in subitizing would be countered by the location and feature confounds. As predicted, subitizing does not occur when displays involve concentric circles (or squares) as stimuli: concentric circles share a center area as well as a center point, which means visual indexing mechanisms struggle to individuate them as distinct objects, and serial attention is required to individuate these (Trick & Pylyshyn 1994, 97). Similarly, when attempting to quantify Os in a sea of Qs, subitizing fails because the target stimuli (the Os) and the distractors (the Qs) share too many features for the targets to be individuated in a way that pop-out perceptually.

To sum up, Pylyshyn's visual indexing theory (Pylyshyn, 1989, 2001, 2007) offers a compelling explanation of how subitizing works that is well supported by behavioral data and consistent with neuroimaging data supporting the two-systems approach. It is important to emphasize that there is no numerical content generated by the visual indexing system responsible for subitizing on this account:

> [Visual Indexing Theory] proposes a limited set of indexes that automatically pick out individual visual objects. The indexing mechanism does not itself encode object properties *nor does it provide a numerical code for the cardinality of the set of indexed items*. It merely provides an indexical reference to the individual objects so that subsequent processes can operate on them. Thus, to derive the cardinality of the set of indexed objects, a subsequent stage of enumeration is required. (Haladjan & Pylyshyn 2010, 308, emphasis mine)

Given that visual indexes have no numerical content, and that my aim is to identify the cognitive systems that allow numerical cognition, I need to discuss how visual indexes figure in a more general enumeration process. Before I can do this, I have to clear up Trick and Pylyshyn's claim that subitizing relies on a preattentive process, given that this claim has come under fire in recent publications that question the validity of this model based on findings that suggest attention is necessary for subitizing.

1.2.6 Attention and subitizing

If Pylyshyn's explanation of subitizing in terms of visual indexing theory is correct, subitizing results from pre-attentive indexing mechanisms, which would seem to imply that variation in attentional load should not affect it. A potential problem with visual indexing as an explanation for subitizing is that several recent behavioral studies appear to have established that, on the contrary, variations in attentional load can significantly compromise subitizing. These findings have been confirmed by brain studies showing that subitizing recruits areas often associated with attention (Burr et al. 2010; Ansari et al. 2007). Studies varying attentional load using dual task procedures (Vetter et al. 2008), the attentional blink (Egeth et al. 2008; Olivers & Watson 2008) and inattentional blindness paradigms (Railo et al. 2008) have all shown that subitizing is compromised when another task recruits too many attentional resources. To see whether or not this affects the worth of Trick and Pylyshyn's account of subitizing, in this section I take a look at some of the studies exploring the relationship between subitizing and attention.

The first study to test preattentive vs attentive models of subitizing used an inattentional blindness paradigm (Mack & Rock 1998) to determine whether variation in attentional resources available for enumeration would have an effect on subitizing (Railo et al. 2008). Inattentional blindness paradigms exploit the fact that we are often blind to stimuli that are not relevant to the task we are trying to complete. In this case, the idea is that if subitizing recruits pre-attentional resources, as Trick and Pylyshyn claim, then it should not be affected even if our attention is focused on another task.

Here, the primary task was to determine which arm of a cross was longer. As a secondary task, subjects were asked to enumerate 1-6 dots that briefly (200ms) appeared outside the area of the target stimulus in non-canonical arrangements. To vary attentional load, three conditions were used. In the inattentional condition, dots were presented without warning after a few trials, at the same time as the primary display. If participants noticed these, they were asked to say how many there were, and to rate confidence in their response on a scale of 1 to 5. In this inattention condition, attentional resources are dedicated to the primary task of comparing line lengths on the cross.[35] In a divided-attention condition, dots would appear only on certain trials, but participants knew this could happen, unlike in the inattention condition. Finally, in the full-attention condition, participants' primary task was to enumerate the dots, even though they were asked to focus gaze on the cross while enumerating.

The underlying rationale here is that if subitizing does not require attentional resources, then determining the quantity of dots should not be affected if their number is in the subitizing range, while performance should be affected outside the subitizing range,

---

[35] Of course, as Railo and colleagues admit, if participants noticed the dots, some attentional resources must still have been taken up by these.

given its reliance on serial attention. Taking three different attention conditions allows researchers to determine whether or not attentional manipulations affect subitizing. According to the authors,

> the preattentive model predicts that enumeration within the subitizing range should already be accurate in the inattention condition and performance within the subitizing range should not be affected by the manipulation of attention, whereas the attentive model predicts that the accuracies should decrease as the number of the objects increases and attention would have an effect even within the subitizing range. (Railo et al. 2008, 87)

Contrary to the preattentive model's predictions, Railo and colleagues found that enumeration performance was affected in both the inattentional and divided-attention conditions, while performance matched that found in regular subitizing studies in the full attention condition. In the inattention condition, accuracy was high only for 1 and 2 dots. Given that this is well below the regular subitizing range of 3-4 objects, the authors conclude that "subitizing cannot be explained by purely preattentive mechanisms" (Railo et al. 2008, 100), since below-par performance was achieved by modifying attentional resources available for enumeration.

While this is certainly plausible, another explanation appears equally possible that does not invalidate the role of preattentive individuation in subitizing. Rather, in line with the finding mentioned above according to which MOT and subitizing compete for indexes (Chesney & Haljian 2011), the reduced subitizing range in this study could reflect the fact that some individuation indexes were taken up by the primary task, meaning that less of these remained to be exploited for fast enumeration. The same interpretation of these findings based on the individuation-hungry demands of attention to a main task could also apply to other recent studies questioning the fact that subitizing relies on preattentive parallel processing.

For example, consider the dual-task procedure employed by Vetter et al. (2008).

Adopting the framework of Load Theory of Attention,[36] according to which attentional demands of a primary task affect processing of a secondary task, they varied the difficulty of a primary task – in this case, color identification – to determine whether variation in attentional load would affect enumeration performance when it is a secondary task. The idea is that if there is a performance difference in the secondary task between low-load and high-load conditions, then that means that attentional resources are shared between the primary and secondary tasks, and thus that these tasks depend on attention. Here, in the low load condition, subjects merely had to determine whether or not a color was present, while in the high load condition, they had to detect two specific color–orientation conjunctions. In both load conditions, the secondary task was to identify the number of targets from a circular arrangement of high-contrast gabor patches ranging from 1 to 8 (see Figure 1.3).

Data obtained show that both load conditions affected subitizing accuracy, with more pronounced effects in the high-load condition. Here too the authors conclude that their results challenge both "the traditionally held notion of subitizing as a pre-attentive, capacity-independent process", and "the proposal that small numerosities are enumerated by a mechanism separate from large numerosities", instead supporting "the idea of a single, attention-demanding enumeration mechanism" (Vetter et al.2008, 1).

However, my reinterpretation of Railo et al.'s (2008) study applies here too: it appears sensible to consider the possibility that the effects on subitizing were due to there being less visual indexes remaining for enumeration due to the demands of the main task. This interpretation explains different effects between high- and low-load conditions:

---

[36] According to load theory of attention (e.g. Lavie & Tsal 1994; Lavie et al. 2004), attentional selection depends on task demands. When a task does not take up too much attentional resources, it is possible to perform another, secondary task, using the leftovers. But when tasks require high attentional load, performance in secondary tasks suffer. There are many ways to vary attentional load. In the present case, researchers claim that if subitizing is the result of preattentive processes, then it should not be affected by the attentional load of a primary task.

the high load condition would recruit more indexes, given that there are more perceptual variables to pay attention to, which would leave less resources for subitizing than in the low load condition.

If I am right, such studies showing attentional effects on subitizing by varying attentional load are not inconsistent with Trick and Pylyshyn's claim that subitizing depends on preattentive individuation mechanisms. The attention required by the primary task can direct individuation in a way that affects individuation in the secondary task even if such individuation occurs pre-attentively. There is ample evidence that the nature of the task can determine at which stage attention is deployed, which means that some tasks can compete for the same pre-attentional resources (Wu 2014, 20).

Of course, even if my reinterpretation hits the mark, the recently available brain data indicating that attentional networks are activated in the subitizing range[37] appear to constitute rather telling proof that attention does affect subitizing. Happily, there is a sense in which it is not debateable that subitizing *does* require attention: subitizing requires a conscious decision, that of enumerating a number of perceptually-presented objects. People do not go about subitizing in the same way they go about breathing or building red blood cells. Given that enumerating is intentional and goal-driven, chances are that subitizing requires attention.[38] This is not a problem for Trick and Pylyshyn, since they do not claim that subitizing itself is preattentive. Their claim is merely that it rests on a parallel individuating mechanism. On the contrary, Trick and Pylyshyn

---

[37] E.g. Pagano et al. 2014; Piazza et al. 2011; Burr et al. 2010; Ansari et al. 2007; Mazza & Caramazza 2015.

[38] While it could be argued that intentional, goal-driven action does not imply attention, this appears unlikely in the case of subitizing. And besides, as Wu points out, goals usually do involve attention: "the subject's goals pervasively influence attention, so much so that some theorists have questioned whether there is attention without the influence of goals" (Wu 2014, 38).

investigated the effect of attention on subitizing and found that cues could improve enumeration even in the subitizing range (Trick and Pylyshyn 1994).[39]

Thus, contrary to the claims of many recent authors discussing the relation between attention and subitizing, the fact that it has been shown that subitizing requires attention is not altogether much of a surprise, nor need it undermine the validity of Trick and Pylyshyn's account, as long as it is possible to explain attentional effects by appealing to preattentive interactions, as I have proposed. So while there is evidence that attention is necessary for subitizing, it seems easy to accept this as confirmation that attentional processing relies on pre-attentional selection, and that tasks compete both for the attentional resources and the pre-attentional selecting and indexes that must accompany any attentional processing.[40]

The fact that neuroimaging data show some kinds of attention are involved in subitizing can actually be considered support for Trick and Pylyshyn's view. For example, as mentioned earlier, evidence was found that subitizing recruits brain areas associated with stimulus-driven attention (Ansari et al. 2007). This accords well with Pylyshyn's visual indexes, which "are assigned primarily in a stimulus-driven manner, so that salient feature characteristics or changes are automatically indexed" (Sears & Pylyshyn 2000, 2). Thus, it is false to talk of challenging "the traditionally held notion of subitizing as a pre-attentive, capacity-independent process" (Vetter et al. 2008, 1), or that "Subitizing was consider [sic] to be a preattentive process for many years" (Gliksman et al. 2016, 1), since the claim has never been that subitizing itself is preattentive, but rather that it *relies* on a particular preattentive visual indexing

---

[39] More proof that cueing modulates attention in the subitizing range can be found in Gliksman et al. 2016.

[40] See Dehaene & Cohen 1994, who noted years ago that attentional processing goes hand in hand with pre-attentional processing.

mechanism.[41]

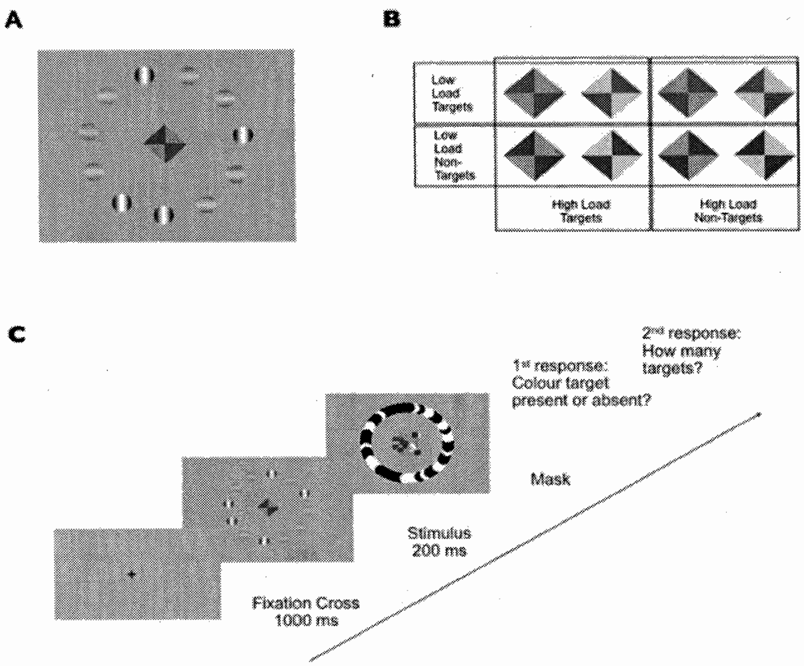

3

Figure 1.3 Experimental set-up from Vetter et al. 2008.

Even if my reinterpretation works, however, the consensus now seems to be that individuation is one of attention's main roles (Cavanaugh 2011). Worse, the very notion of a distinction between preattentive and attentive processing has been

---

[41] The same reasoning applies to the claim that object selection in MOT could be directed by attentional processing (Oksama & Hyona 2004): participants in MOT are asked to pay attention to objects and track them, so there is clearly an attentional component here too. But this attentional component's essential contribution to these processes must not be overstated: if I intend to take a deep breath and then attend to my breathing, does that make breathing an attention-dependent process?

questioned, as has the possibility of preattentive processing (Joseph et al.1997; Duncan et al.1989). This means that it may be impossible for Trick and Pylyshyn's visual indexes to be accurately described in terms of preattentive individuation processes, since there may be no such thing.[42]

At this point, it is evident that the investigation on the nature of subitizing has brought us at the heart of a complicated question, that of the relationship between attentive and preattentive processes, and how objects are individuated. This topic has important implications for a number of philosophical issues. For example, as Spelke puts it, "[t]he parsing of the world into things may point to the essence of thought and to its essential distinction from perception" (Spelke 1988, 229).

Unfortunately, as has often been remarked, attention is a particularly confused, yet important, notion (Mole 2014; Wu 2014), and a number of important questions in attention research have not found solutions because "the key concepts (selection, automaticity, attention, capacity, etc.) have remained hopelessly ill-defined and/or subject to divergent interpretations" (Allport 1993, 188, cited in Wu 2014). Thankfully, our purpose in this chapter is not to settle these matters, but 'merely' to identify the cognitive systems underlying enumeration, so as to make our way towards the origin of numerical cognition. In this context, the precise role of attention in subitizing need not be settled immediately. This is because what matters is whether or not Trick and Pylyshyn's individuating mechanism allows us to explain behavioral discontinuities in

---

[42] As Vetter et al. put it:

> the strong notion of pre-attentive/attentive dichotomy has been regarded as an oversimplified account in the attention literature...and particularly the hypothesis of attention-free perceptual processing has been questioned. Indeed there is evidence that even the simplest forms of feature detection (e.g. orientation detection), which had previously been thought of as occurring pre-attentively, depend on the availability of attentional resources in a dual-task situation. (Vetter et al. 2008, 1)

enumeration tasks, regardless of whether or not it takes place pre-attentively.[43]

For now, let me simply point out that there are many types and levels of attention (Wu 2014) so even if both counting and subitizing require attention this does not mean they recruit the same cognitive systems, since they can have different functions within an attentional processing hierarchy. For example, even if individuation requires attention, it can still occur in parallel, while this is not true of counting. So even if it does turn out that attention is required for subitizing, the (now-attentive-instead-of-preattentive) individuation mechanism plays the same part in subitizing that it does if individuation is preattentive – namely, selecting feature clusters in parallel to allow (higher-level) attentional processing on these. The important insight that Trick and Pylyshyn offer is that subitizing relies on an individuation mechanism that tags a limited number of objects for further processing:

> A critical aspect of the explanation is the assumption that individuation is a distinct (and automatic) stage in early vision and that when the conditions for automatic individuation are not met, then a number of other phenomena that depend on it, such as subitizing, are not observed. (Pylyshyn 2003, 175)

In sum, while there remains a number of controversies surrounding our ability to enumerate, including the number of cognitive systems involved and the role of attention, Trick and Pylyshyn's visual indexing account still affords a plausible and well-supported explanation. Having argued for the merits of the visual indexing account in this section, in the next section, I explore what role visual indexes could possibly play in enumeration.

---

[43] Illustrating that progress can be made without settling the preattentive-attentive divide, some authors (Pagano et al. 2014; Mazza 2017) simply speak of an attention-based individuation mechanism behind subitizing.

1.3 The Object-File System

1.3.1 From visual indexes to object files

I have just argued in defense of Trick and Pylyshyn's account of subitizing, according to which the difference between enumerating inside and outside the subitizing range can be explained by the accelerating effects of a limited-capacity parallel individuation system. While I have shown that support for their account is strong, important details concerning how this system manages to contribute anything abstract enough to figure into enumeration processes are lacking. Recall: Trick and Pylyshyn's two-stage description of subitizing involves an initial parallel individuation stage with no numerical content followed by "matching each individuated item with a number name, in the order of the number name" (Trick & Pylyshyn 1994, 88).

Given that visual indexes are, well, visual, and that numbers are not, we need to know more about the way in which information from visual indexes becomes available for more general, amodal processes like enumeration. After all, visual indexes are made up of clusters of visual features like line orientations and color changes. On the other hand, enumeration involves tagging objects presented in many modalities using abstract representations with numerical content. So the question is, if subitizing relies on visual indexing, how does visual indexing allow 'matching' indexes to number words in enumeration?

To see how visual indexes can figure into more general processing, it is essential to understand what sort of information about the world they carry. First, note that to enumerate visually perceived objects - or undertake any other action involving visual stimuli, for that matter - these must first be individuated as particular objects, separate from the rest of the visual field. There are many levels at which the visual system parses incoming sensory information to yield representations of particular objects. For example, discontinuities in features of the visual field (e.g. colors variations, brightness levels, line orientations and curvatures, etc.) can be grouped together to form object boundaries. Features are detected in parallel throughout the visual field, and then grouped together by similarity and proximity to form indexes of feature clusters (Marr 1982).

Such indexing cannot be based on individual features of the objects, since these are typically in constant flux. For example, imagine watching a car drive by on a partly cloudy day. While our experience of the event is of one spatiotemporally identical object passing in front of us, in actuality, the car's color and brightness levels vary as it passes through different patches of sunlight, and its retinal size fluctuates according to its distance from us. Given such variation, it would be impossible to keep track of the car based only on its retinal features. To accomplish tasks like enumeration of visually presented objects, what we need to keep track of objects is a way to individuate objects despite variations in their features and position. Visual indexing mechanisms allow such reference tokening to take place by tracking clusters of features as they move through the visual field. To highlight the fact that these individuation indexes ostensibly point to feature clusters like a finger points to an object, Pylyshyn dubbed them FINSTS, for FINgers of INSTantiation.

Seen this way, visual indexes are referential: they point to things in the world (Pylyshyn 2007). And yet, as mentioned in the MOT discussion, they are limited to tracking spatiotemporal continuity and do not carry feature information along with their referent. Thus, visual indexes are limited to the demonstrative content THIS, individuated about four times in parallel. Visual indexes, then, are part of a mechanism for demonstrative thought that participates in representation, but is not itself representational, since visual indexes lock onto objects, but not their properties.

So the question is: given their purely demonstrative role, how do finsts allow us to represent anything about the world? When a finst is individuated, the properties of the objects that triggered this individuation are not encoded along with it. For such properties to be attached to the object that triggered the opening of a visual index, a further processing step is required: the opening of an object file (Kahneman et al. 1992).[44] Finsts link objects in the world to object files, which can represent properties of objects that trigger individuation of visual indexes and make this information available for further computations. Thus, the function of visual indexes is to keep track of objects and allow us to attribute properties to these via object files.

A good illustration of what sort of content object files can carry can be found in the first study to use the term 'object files' (Kahneman, et al. 1992), which used an object-reviewing paradigm to study potential facilitation effects of object individuation on feature recognition. Here, the setup involved exposing subjects to screens on which two identical squares were displayed, one on top of the other. In each square, a letter was then briefly presented and hidden from view. The squares then moved to new locations, equidistant from their initial position, after which a letter appeared in one of them. Subjects were tasked with identifying this second letter as quickly as they could. The objective was to determine whether the initial letter, despite being displayed in a different location from the second, could prime (i.e. facilitate) semantically-related responses in a different location. This was indeed the case, which suggests that content-based effects like semantic priming follow the same course as the tracked squares, despite the fact that the priming letters had been displayed in separate locations. Kahneman and colleagues took this as evidence of a mid-level individuating mechanism that allows us to track objects' motion as well as some of their features and stores this information in object files.

Given that both object files and finsts are limited in number and track individuals, it is not surprising to learn that object files and visual indexes share many properties:

---

[44] See Mitroff et al. 2005 for a discussion of object files and their relation to conscious percepts.

> Visual indexes point to features or objects, not to the locations that these stimuli occupy. Like the object files of Kahneman, Treisman, and Gibbs (1992), visual indexes are object-centered and continue to reference objects despite changes in their location. (Sears & Pylyshyn 2000, 2)

In fact, depending on who one reads, the two notions are equivalent, as illustrated by some authors having described subitizing as "the rapid enumeration of open object files, without counting (verbal or nonverbal)" (Cordes & Gelman 2005, 138).

While the relationship between finsts and object-files is, to use Pylyshyn's words, "not entirely transparent" (2003, 209f), he does mention that Kahneman and colleagues' object files and his finst-based theory of visual indexes were developed for different purposes: while "object file theory has emphasized memory organization and its relation to the objects from which the information originated", his visual index theory "emphasizes the mechanism required for establishing and maintaining the connection between objects in the visual field and mental constructs (representations) of them" (Pylyshyn 2003, 209f).

We see that visual indexing theory focuses on mechanisms that 'connect' feature clusters in the visual field and the objects themselves, while object files are the equivalent of memory slots where information about objects is stored (see Figure 1.5). Importantly, since object files refer to the world via visual indexes, they share the same capacity limits: only a small number of object files can be individuated at once. The difference is between recognition and individuation. Only object files can be considered as representational.[45]

---

[45] See chapter 4 of Pylyshyn 2003 for more details on this distinction, as well as Pylyshyn 2007, 37-39. Further, note that there is some evidence that the ability to index objects may come before the ability to store object properties in objects files (Leslie et al. 1998; see also Wu 2014). There is also evidence that not all features can be clustered at the same stage in development (Kàldy & Leslie 2003)

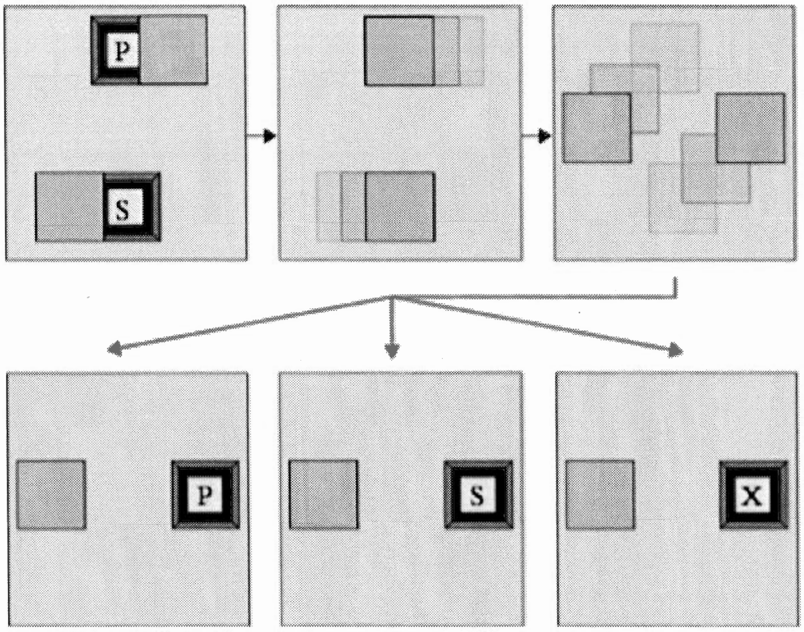

4

Figure 1.4 Experimental design in Kahnemean et al.1992

Happily, despite potential ambiguity in separating finsts and object files, what is important for us here is that there is ample evidence that a system capable of attaching properties to a limited number of objects in parallel can explain enumeration in the subitizing range. This section has shown that even though it is the finsts that individuate the perceptual data into visual indexes, it is the object files that carry the information to other cognitive systems. Since it is the object files that represent properties of the world, it is the object-file system that can figure into an enumeration process, where number words can be assigned to individuated objects. This means that to find out the origin of numerical cognition, we need to know where we get our object files from. In the next few sections, I discuss evidence for the existence of an object file system in human infants and nonhuman animals, indicating that the object-file system is innate.

Along the way I describe some of the research methods used so as to highlight what sort of information we can take from data obtained using these methods, and to determine how the object-file system could support numerical behavior.

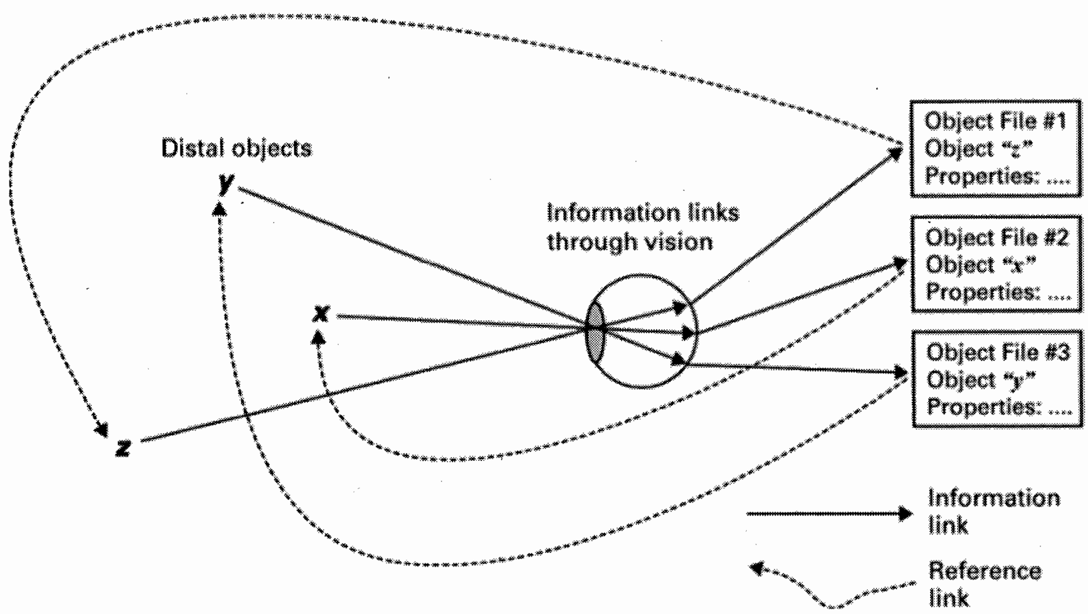

Figure 1.5 Relation between object-files and finsts

1.3.2 Evidence of innateness and phylogenetic origins

I have just discussed how object files allow us to store content and apply it to visual indexes. To find the origins of the systems deployed in enumeration tasks – at least, in the subitizing range – we need to know where object files come from. Given that, in the visual domain, object files relate to the world via pre-attentive visual indexing devices, it is not unreasonable to expect to find evidence of their presence in preverbal infants and nonhuman animals. Also, if we accept that object files relate to the world via finsts, we can expect them to share many of their properties – in particular, capacity limits of about 4 objects, primacy of spatio-temporal information, and distinction of occlusion from cessation of existence. In this section, I take a closer look at the object file system, its origins, and the methods used to determine these. This will allow me to show that it does not have any numerical content, despite being able to support operations about such numerical content, in section 1.4.

Searching for the origins of numerical cognition forces us to consider a number far-reaching philosophical issues. In the case at hand here, the question we are interested in is where our ability to aprse the world into objects comes from MAYBE THAT SPELKE CONTINUOUS TO SPECIFIC QUOTE HERE?. The question of how we individuate objects as materially extended mind-independent entities that persist through time has been around since at least the heyday of British empiricism and its reaction to Descartes and other nativists. The empiricists, on the one hand, allowed very few items in their inventory of innate faculties, limiting themselves to what the senses could give us. This forced empiricist accounts of object individuation to appeal to general learning mechanisms and statistical generalization operating over sensory and perceptual representations to explain how concepts such as OBJECT get their properties (e.g. solidity, spatiotemporal persistence, impossibility of passing through other objects, etc.) In this case, the learning would be piecemeal and gradual, and dependent on the amount of stimulus presented to the learner.

If we opt for a nativist explanation of the origin of representational systems, on the other hand, we should expect to see the effects of the natural maturation of these, such as regular behavioral milestones (think of babies learning to walk and talk). While philosophers could argue all they wanted about where such general concepts came from, little data was available to settle the issue.

In a rare and wonderful example of science coming to settle philosophical matters, data collected by recently developed experimental paradigms have allowed researchers to probe deeper into the cognitive machinery of neonates, preverbal infants, and nonhuman animals. In recent times, experimental methods have been developed that seem to have settled the matter rather conclusively, with nativist accounts of the origin of object concepts gaining the upper hand (Carey 2009). In what follows, I discuss the main empirical findings supporting the existence of an innate, phylogenetically-inherited system representing object files. First, I sketch the thinking behind the main source of evidence for the presence of this system in infants and animals, Violation-Of-Expectancy (VOE hereafter).

1.3.3 Violation of Expectancy

A difficult hurdle when probing infant and animal cognitive systems is that they cannot verbally report on the causes of their behavior. Thankfully, in recent years, methodological breakthroughs appear to have opened the door into their representational repertoire. Techniques such as Violation-of-Expectancy (VOE), habituation, and manual search have allowed researchers to detect the presence of representations in subjects that were out of reach when Piaget first started looking into the minds of infants.

Consider the VOE paradigm, which in recent years has become the most widespread methodological tool used to investigate the representational repertoire of infants and animals (Carey 2009). VOE is a specific application of looking time studies, which were initially geared towards detecting perceptual skills in infants (e.g. whether or not they could see colors). In VOE, researchers exploit subjects' tendency to look longer at stimuli that presents novel or unexpected scenes. The underlying rationale is that we can detect the presence of representational systems geared towards certain aspects of the world by observing whether or not violations of this aspect of the world incur longer looking times in the exposed subjects. In other words, subjects would not stare longer at outcomes that violate the way they expect the world to behave unless they have a representation of this aspect of the world.

Many VOE experiments involve little – if any – training, simply exposing subjects to unexpected scenes and reading their reaction. Other VOE paradigms involve repeated exposure to perceptually similar scenes. In these habituation paradigms, researchers exploit the fact that subjects eventually lose interest (i.e. habituate) when repeatedly exposed to stimuli displaying the same properties. Typically, after a series of trials that fail to present novel features, subjects' gaze starts to wander away from the stimulus, which appears to be presenting the same content over and over again.

As an indicator of habituation, researchers typically set a value that is proportional to the looking time for the first stimulus. For example, if subjects first looked at the stimulus for four seconds, when they start to look at stimuli for 2 seconds before looking away, it is taken as a sign that the stimuli are no longer presenting anything new or interesting to the subject.[46] Once subjects have habituated, researchers expose them to stimuli that are meant to regain their interest, either by showing them possible modifications of previously displayed stimuli, or modifications that, to an adult, would be impossible or unexpected. The important behavioral datum here is whether looking time varies between habituation stimuli and novel stimuli.

Here the idea is that if repeated exposure to the same stimulus eventually bores subjects, when a change in certain features of the stimulus is able to get looking time back up, it is interpreted as a sign that the subjects' reaction is based on a representation that is sensitive to those changed features. In particular, stimuli that display impossible outcomes would elicit longer looking times given that they violate the expected outcome as determined by the constraints on the representational systems responsible

---

[46] In some experiments, variations in looking time are relatively small (around two seconds) (e.g. Xu & Carey 1996), but in all cases, they are reliable and replicable differences in behavior, and while explanations may vary with respect to the richness of the representations used to explain these behavioral differences, all seem to force some kind of innate representational content.

for subject's reaction to the stimuli.[47]

As an example, consider the pioneering work of Renée Baillargeon and her colleagues (Baillargeon et al. 1985), who were the first to test the presence of representations of object permanence using a VOE paradigm that involved playing a trick on the infant subjects. In this study, four-month old infants were shown a platform that moved back and forth 180 degrees, like a drawbridge that falls in two directions. Once habituated to the motion of this platform, infants were shown an object introduced directly in the path of its downward motion. Normally, this would mean that the platform could not complete a full 180-degree rotation, since the object had been placed in a way such as to prevent this motion.

In some trials, infants were shown the expected outcome, and the platform halted when it touched the occluding object before going back in the other direction. In other trials, an impossible outcome was shown to the infants: without the infants noticing, researchers surreptitiously removed the object from where infants had seen it placed. Thus, instead of stopping on the object, the platform continued its downward motion as it had done when there was no occluding object. Infants looked longer at the impossible outcome, thus suggesting that they expect objects to be solid to the extent that one cannot pass through the other. Importantly, the fact that infants looked longer at the impossible outcome suggests that the representational system underlying their behavior in this setup can keep track of objects even when they are not being directly perceived, which suggests that infants are equipped with an innate representation of object permanence.[48]

---

[47] Experiments such as Wynn's arithmetic tasks described below and its replications depend on the violation of expectation paradigm, which assumes that longer looking times indicate an infant's violation of expectation. However, some developmental psychologists (e.g. L. B. Cohen & Marks, 2002; Haith 1998) challenge this core assumption, and account for the results in purely perceptual terms.

[48] As Carey (2009) convincingly argues, this shows that both Piaget and Quine were wrong in thinking

A great advantage of the VOE paradigm is that it allows researchers to probe deep into the representational inventory of subjects that cannot produce introspective reports - especially, preverbal infants and animals – using a non-invasive behavioral reading. If subjects react in ways that display the presence of underlying representations with the same features as those of human adults, VOE may point to the presence of ontogenetic or phylogenetic continuity in underlying representations used for certain cognitive or perceptual tasks. We can thus piece together the contents of the representations behind infant and animal behavior in various tasks by taking a look at the inferences that infants appear to be making concerning possible and impossible outcomes. In the following section I review evidence showing the presence of an object file system in human infants and animals that shares its three processing signatures mentioned above.

1.3.4 Three properties of object files in infants and animals

I showed above that visual indexes prioritized spatiotemporal information over other features, had a limited capacity of about 4 objects, and distinguished between cessation of existence and occlusion. I also mentioned that the OFS shared these characteristics. In this section, I show that these three characteristics also appear to underlie infant and animal object representations, thereby showing that our object file system is both innate and phylogenetically inherited. I start with occlusion versus cessation of existence.

In the first experiment to use habituations methods (Kellman & Spelke 1983), infants were exposed to a screen showing a moving partially occluded rod. Then, the occlusion was removed to reveal either a solid moving rod (expected) or two moving rods which lined up with each other but had a gap where the occluder had been, thereby creating

---

that infants had to learn about object permanence.

the illusion of a single moving rod (unexpected) (see figure 1.6). Infants looked longer at the unexpected outcome, and this was interpreted as a sign that infant perception displays certain gestalt principles, including amodal completion of partially occluded objects. Replications of this behavior are taken as evidence that infants represent objects existence as persisting despite the fact that they are not being perceived, either due to partial or total occlusion.[49]

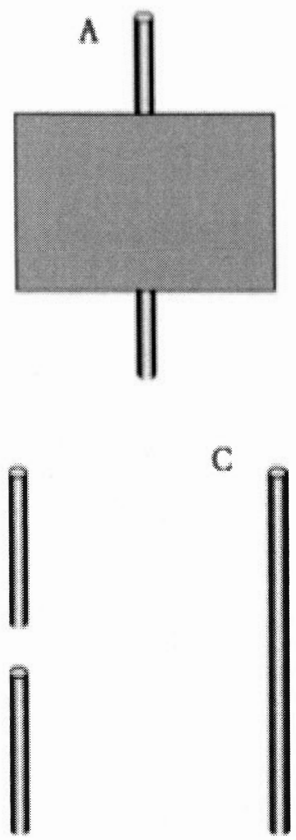

Figure 1.6 Schematic depiction of occluded rods in Kellman & Spelke 1983.

[49] See Valenza et al. 2006, who used stroboscopic presentation of similar stimuli to probe the existence of amodal completion of object contours in neonates. See also Johnson 2010.

Given the widespread use of MOT to study object individuation in adults, it is not surprising that the same methods were used to study object files in infants. For example, Cheries et al. (2005) habituated infants to displays showing disks moving randomly, sometimes running into occluders (bars). As in adult studies, the disks then either appeared to pass behind the occluder, or to disappear on one side before reappearing on the other. Then, in target trials, the number of moving disks changed. Only those infants who had been shown movement consistent with real-life object occlusion dishabituated to the novel number. This suggests that those shown objects disappearing and reappearing lost the ability to track their number. Such findings appear to confirm the same distinction in infants between objects' ceasing to exist and objects being occluded from sight seen in adults.

Similar experiments with newly hatched chicks show that these group together around a fully formed triangular-shaped object after exposure to a partially occluded triangle, which suggests the presence of an innate system that, when confronted with partially occluded objects, individuates these by filling in the occluded contours (Regolin & Vollortigara 1995). Similarly, chicks will search for objects on which they have imprinted when these are hidden from view despite no previous exposure to any objects whatsoever, given that they were just hatched. This suggests that instead of computing object permanence from statistical learning over sense data, chicks are equipped with perceptual input analysers sensitive to the spatio-temporal continuity of object permanence. Such results suggest that chicks are equipped with an innate capacity to represent the spatio-temporal continuity of object permanence despite the fact that objects are not directly being perceived.[50]

---

[50] Representational systems displaying object persistence despite occlusion have also been demonstrated in chimpanzees (Beran 2004), rhesus macaques (Hauser & Carey 2003), and tamarins (Uller et al. 2001).

1.3.5 Capacity limits

Many experiments with infants and animals have displayed the capacity limits of MOT. Among the earliest to show such capacity limits in infants are Starkey & Cooper (1980), who habituated 4- to 7-month-old infants to displays of two dots. Their looking time increased when they were shown displays of three dots. Symmetrical results were obtained for infants habituated to 3 dots. However, when the variations involved 4 vs 6 dots, infants failed to dishabituate. The fact that infants failed to discriminate between displays of 4 vs 6 dots but not 2 vs 3 dots despite the same ratio is evidence that the system underlying their abilities here is limited to about 4 objects, thus mirroring capacity limits in adult object files.

One animal species that does display the same capacity limits to their object representations are rhesus macaques. Hauser and colleagues (e.g., Hauser & Carey 1998, 2003; Hauser et al. 2000) were the first to apply VOE and manual search paradigms to both free ranging and laboratory-housed primates to probe their object representations and find data that converges with that obtained in infant studies. Adopting Wynn's (1992a) VOE paradigm (see below) to free-ranging rhesus macaques, Hauser and colleagues used disappearing tricks to expose primates to stimuli that violated object permanence in order to test their capacity to perform basic numerical operations like addition and subtraction. Here, in impossible displays, subjects were shown two eggplants placed in a display box behind a small screen. When the screen went up, they could only see one eggplant, the other having been secretly hidden in a pouch. Converging with infant studies, the results showed that subjects looked reliably longer at impossible outcomes than at possible outcomes.

Similar methods were used to show that macaques look reliably longer when shown impossible outcomes of tasks that some authors describe as instances of subtraction. Here, two eggplants are inserted on a display, a screen goes in front of them, they can see two eggplants being removed, and then when the screen is lowered they find the impossible outcome that there is still an eggplant on the display. Importantly, later experiments showed that these abilities break down when tasks involve four objects, since they fail to look longer at impossible outcomes such as 2+2 = 3, despite looking longer at tasks displaying 2+1 = 2. In another paradigm, macaques were shown apple slices inserted into a bucket, followed by another number of apple slices into another bucket. They were then allowed to choose one of the buckets. They reliably chose the larger number of apple slices when buckets contained 1 vs 2, 2 vs 3 and 3 vs 4 slices, but could not perform well with 2 vs 5, 4 vs 8, and even 3 vs 8 slices, despite the large discrepancy. This suggests that the same capacity limits on the underlying representational system as present in human object files (Hauser & Carey 2003).

1.3.6 Primacy of spatio-temporal information in object files

As for the primacy of spatio-temporal information in object files, Richardson & Kirkham (2004) adapted paradigm used by Kahneman and colleagues in their study of object files (described above in section 1.3.1) and exposed 7-month old infants to the same visual presentation of boxes arranged vertically. Infants then saw objects appear and disappear in the boxes. For example, in the top box, a duck would appear, and a quack noise was played before the duck disappeared from the box. In the bottom box, a bell appeared and a ringing sound was played before the box disappeared. Then, the empty boxes moved to a new location, each equidistant from each other and from their previous location, and a sound was played. The question was: where would the child look?

If infants individuate objects according to spatio-temporal continuity, then they should look to the box where the object associated with the sound was initially displayed, since that box would have the sound as part of its attached features. This is indeed what researchers found, suggesting that infants' object files, like those of adults, individuate objects based on the spatio-temporal continuity of feature clusters.[51]

As for animals, in a clever experimental setup, researchers found a way to determine whether primates could also individuate objects by prioritizing spatiotemporal information (Flombaum et al. 2004). Rhesus macaques observed as experimenters rolled lemons down a ramp until they stopped behind an occluding tunnel. Out of the tunnel came a kiwi, which was then hidden from view behind a second occluder. If the motion of the kiwi was continuous with that of the lemon as it passed through the tunnel, the macaques only searched for food in the second occluder, which suggests that they individuated the food not on its individual features, but on its spatiotemporal features. On the other hand, if the lemon stopped long enough for the motion of the kiwi to appear discontinuous, the animals searched for two morcels of food. This behavior suggests that these monkeys prioritize spatiotemporal cues when tracking moving objects, rather than features (e.g. those that distinguish lemons from kiwis), like human infants and adults have been shown to do.

1.4 Does the Object-File System have numerical content?

I have just presented evidence showing that the three signature properties of the object-file system (OFS) – a capacity limit of about four objects, priority of spatio-temporal information, and distinction of occlusion from cessation of existence – have been found in human infants, adults, and nonhuman animals. Given that the limits on the number of objects we can enumerate quickly and effortlessly is the same as the limit to the number of objects that infants and animals can simultaneously track, it looks like the

---

[51]There are certain conflicting results in the literature concerning when infant behavior displays the use of spatio-temporally continuous features to individuate objects, rather than, say, kind information, but converging evidence from manual search studies confirms that infants below 12 months are more sensitive to the former (see Xu & Carey 2000; Carey 2009: chapter three).

system responsible for our ability to subitize is also present in infant and animals.

However, so far none of these features appear to be related to numerical abilities. On the contrary, the fact that we can only open about four object files at the same time would appear to make it an unlikely candidate for the origins of numerical cognition, given that there is no such limit to numbers. And yet, the reason we are currently investigating the object file system is that we are trying to track down the system that explains behavioral discontinuities in the subitizing range, and understand how it could be involved in subitizing. Now that we know more about the OFS, it's time to ask: does it have any numerical content?

Recall that I mentioned above that Pylyshyn claims that performance in the subitizing range can be explained by the properties of a parallel individuation system that does not have any numerical content, since visual indexes to not attach any features to the objects they track. Rather, the visual indexing system's job is merely to provide demonstrative tags for further processing, so that their content would be limited to something like THIS, THIS, THIS, for three objects. I also mentioned that, on Pylyshyn's account, the numerical aspect of subitizing comes in *after* visual indexes have been attributed to objects, when the output of a visual indexing system is matched to a memory slot for number words. In subitizing, then, the visual indexes are limited to individuating objects, while a separate process reads numerical content from the individuated objects and matches it to a numerical representation. But, given that object files, unlike visual indexes, can encode properties of objects, it is not unreasonable to wonder whether the OFS can represent numerical information about the world on the basis of relations between the indexes it operates over.

Support for this possibility comes from infant and animal behavior in the subitizing range that is often taken as evidence of elementary arithmetical skills (Dehaene 1997/2011; Wynn 1992a; Carey 2009). Scores of studies in the past few decades have found that infants can discriminate between small numbers of objects based on their numerosity and appear have expectations about the arithmetical relations between these objects. The question is whether the OFS, which allows features to be bound to objects, is responsible for this behavior, and whether or not this means that this system computes numerical content.

Consider for example the most famous study of infant numerical skills. Here, Karen Wynn (1992a) used a VOE paradigm to test the number of hidden objects infants (aged 5 months in some trials) could keep track of. To do this, Wynn set up a small stage on which one or two objects (depending on the trial) were introduced in sight of the infants. Then a screen went up, and infants saw a hand place another object behind the screen (in the 'addition' condition), or take away an object from behind the screen (in the 'subtraction' condition). The screen then came down to reveal either the right number of objects (possible outcome) or, when experimenters played tricks on them, the wrong number of objects (impossible outcome).

The question was whether infants would expect the right number of objects to be revealed when the screen came down. Wynn's results show that infants look longer at unexpected outcomes. For example, they show no signs of surprise for the correct outcome of simple addition and subtraction tasks like 1+1 = 2, or 2 - 1 = 1, but they stare longer when an object has been surreptitiously added or removed from behind the display, instantiating 1+1 = 1 or 2 - 1 = 2, or 1+ 1 = 3. Showing the capacity limits of the object file system, infants did not show preferential looking times when exposed to 3 vs 4 objects, nor 3 vs 8 objects, suggesting that they lose track when too many objects are presented. Wynn's (1992a, 1998) interpretation of her findings is that infants' reactions are based on rudimentary arithmetical skills, since they react to violations of expectancy that can be described in arithmetical terms. She takes this as evidence that the OFS is capable of supporting rudimentary arithmetical operations, which would support attributing some form of numerical content to this system.

While common, such a rich interpretation of these findings remains controversial.

Instead of attributing rudimentary arithmetical abilities to infants based on their reaction to these impossible events, it would appear more parsimonious to explain their behavior in terms of non-numerical content. There are at least two other plausible interpretations of these findings that do not appeal to any numerical content to explain infant behavior. On the one hand, infants' reactions could be due to violations of their representation of continuous magnitudes that co-vary with numerosity, such as total surface area, contour length, convex hull, or average size. On the other, infants could be reacting to a mismatch between states of the OFS without explicitly representing the reason for this mismatch. I explore these possibilities in turn.

To test for this, a number of follow-up studies varied non-numerical cues in similar transformation paradigms to determine whether infant behavior resulted from a representation of numerical content rather than other perceptual cues. While Simon's (1995) study replicated Wynn's findings and showed that infants are not reacting to variations in object identity, other subsequent studies found that when non-discrete magnitudes like total surface area or total contour length (i.e., the sum of the contours of individual objects) in an array were controlled for, infants reacted to changes in these, and did not react to numerosity changes (Clearfield & Mix 1999; Feigenson et al. 2002).[52]

While these findings invalidate Wynn's arithmetical interpretation of her data, other experiments involving cross-modal matching tasks appear more congenial to a numerical interpretation of infant behavior. Unlike Wynn's VOE transformation study, where stimuli were presented in a single modality, Starkey et al. (1990) tested whether

---

[52] See also Uller et al. (1999) for an excellent summary of the pros and cons of the object file model vs Wynn's arithmetical explanation. Note that while Feigenson and colleagues question the fact that infants react to variation in numerical content in the subitizing range, they do consider some studies (e.g. Xu & Spelke 2000; Xu 2000, discussed below) as providing "unequivocal evidence for sensitivity to number" (Feigenson et al. 2002, 36) above the subitizing range.

infants can represent numerical correspondences between displays of objects presented visually and sequences of sounds, using an auditory-visual preference procedure (Spelke 1976). They presented 6- to 8-month old infants displays containing 2 or 3 objects of various types, meanwhile playing sound sequences of either two or three drumbeats. They found that infants looked longer at displays that were numerically equivalent to auditory sequences, even when duration was equal for both auditory sequences, suggesting a preference for intermodal correspondence based on their numerosity, even in the subitizing range.

Behavior in this experiment cannot be explained by non-numerical cues in a single modality, as was done for Wynn's studies, since behavioral discontinuities are between two cross-modal conditions, which suggests that the underlying representation must not be modality specific, as is the case for total contour length, for example. Thus, this experiment appears to provide evidence for a modality-independent, amodal representation that is sensitive to numerical correspondence between stimuli across modalities.[53]

And yet, despite the fact that these findings appear to eliminate the influence of non-numerical perceptual magnitudes on behavior, this does not mean that infants reactions are based on numerical content. While this type of study is often considered strong evidence for the presence of numerical content in the subitizing range, it is important to consider that infant reaction here can simply be a matter of establishing one-to-one correspondence between object files in both modalities, and finding a mismatch. On this alternative interpretation, infants are reacting to violations of discrete quantities of objects, not to any explicit representations of numbers. Consider how this might apply

---

[53] It is worth mentioning that some experiments failed to replicate these findings (Mix et al. 1997), though later experiments using similar methods appear to show preferences for cross-modal matching in 7-month olds, using more natural stimuli: faces and voices (Jordan & Brannon 2006). See also Izard et al. 2009.

to Wynn's experiment above: here the OFS would produce outputs like OBJECTa and OBJECTb when two objects are shown on a stage. Then, when a screen hides these objects and a third object is added to the stage, another object file, OBJECTc would have tracked it as it was placed behind the screen. Thus, when the screen is removed to reveal the objects it was hiding, infants would expect OBJECTa OBJECTb OBJECTc, say. Due to the experimenter's tricks, however, they perceive OBJECTa OBJECTb only, and this does not match the output of the OFS. On this model, the OFS could allow computations based on a SAME/DIFFERENT distinction, without there representing any explicit numerical content.

A similar interpretation could also explain behavior in Starkey et al.'s (1990) study. Here, infants could open object files in which visual and audio information are stored. Given that matching displays and sounds are easier to track than combinations of 3 visual stimuli and 2 audio stimuli, which involves a missing audio stimulus, infants could prefer the simpler stimuli with matching numerosities, without necessarily representing their numerosity.

At issue here is a crucial question in the investigation of the origins of numerical cognition, that of the numerical content of the OFS. As I have shown, for now, there is no conclusive evidence that this system provides us with anything like the content NUMBER. More parsimonious, non-numerical explanations can account for all the data giving the appearance of numerical abilities in the subitizing range by appealing to objects files, whose existence is well supported. On these non-numerical explanations, the OFS can only be described as supporting one-to-one correspondence between individuated object files, which can lead to representations with the content SAME or DIFFERENT. The underlying reason for this difference or sameness, numerical inequality, does not need to be explicitly represented by infants and animals in this range. In sum, while there is evidence that some behavior in infants in the subitizing range is based on numerical matching between senses, this does not mean that the OFS provides numerical content.[54]

Of course, while subitizing may be an important phenomenon to study the origins of our numerical abilities, the range in which it operates is nothing to write home about: at around 4 objects, the OFS, on its own, can only account for numerical abilities in a very restricted range. We are then left to wonder what system underlies numerical abilities beyond the subitizing range. In the next chapter I present data supporting a representational system with content that can be described as numerical, though even here there are limitations that prevent us from attributing the content NUMBER to this system.

---

[54] Cheung & LeCorre (2018) provide evidence that children as young as two years of age can use the OFS to compare collections based on their numerosity in some cases, but this does not mean that younger infants have this ability, nor that the conditions used in their study single out numerosity as the sole basis for behavior, as discussed below in section 2.5.

CHAPTER II

THE APPROXIMATE NUMBER SYSTEM

2.1 Introduction

We just saw that our ability to subitize is restricted to very small quantities of objects. Of course, the subitizing range's limited domain does not prevent us from representing large quantities of objects in our environment. Most adults regularly and effortlessly represent the number of objects in an attended perceptual space – its numerosity – even though this space contains much more than three or four objects. Since the end of the 20th century, hundreds of studies targeting numerical skills in adult humans and infants, as well as many animal species, have gathered mountains of evidence supporting the existence of a representational system that allows individuals to represent quantities of objects in their environment, with decreasing precision as quantities increase. This so-called 'approximate number sense' (hereafter, ANS) plays a central explanatory role in most recent accounts of the origins of numerical cognition.[55]

For many, this numerosity-dedicated system contains an amodal representation of

---

[55] Dehaene 1997/2011; Feigenson et al. 2004; Gallistel & Gelman 1992, De Cruz 2008; Menary 2015a; though see Gebuis et al. 2016 for a critical review, and see Leibovich et al. 2017 and Lourenko 2015 for alternative accounts of the system underlying numerical abilities outside the subitizing range, discussed below in section 2.5.

numerical quantity that serves as the foundation to our arithmetic and mathematical abilities (Butterworth, 2005; Dehaene, 2009; Piazza, 2010). Given its central role in the development of our arithmetical skills, it is crucial to understand the methods used to establish the existence of such a system as well as the evidence suggesting that its content is numerical. In this section I summarize a few key studies in the short history of the ANS that are taken as support for the existence of such a numerical system in adult humans (2.2), infants (2.3), and animals (2.4). I then briefly summarize the main lines of a recent alternative model of the ANS, the Analog Magnitude System (AMS hereafter), explain the critical motivations for this view, and what this means for my quest for the origins of advanced numerical cognition (2.5).

2.2 Support for the ANS in adult humans

2.2.1 Studying numerical abilities with symbols

As was the case for subitizing and object-file studies, studies probing the ANS proceed by having subjects perform numerical tasks and then looking for performance signatures in the behavioral data obtained. Such performance signatures can reveal the distinguishing properties of the representational system underlying behavior in numerical tasks, thereby shining light on the origins of our numerical abilities. The main experimental paradigms used in the study of numerical skills outside the subitizing range in adults involve numerical identification, estimation and comparison tasks using non-symbolic (e.g. dot arrays, sequences of tones), as well as symbolic (e.g. Indo-Arabic numerals and number words) stimuli, in various modalities.

Perhaps the most famous study of numerical skills outside the subitizing range is also the one that is considered the starting point for the modern study of numerical cognition. This is Moyer and Landauer's (1967) study, which was the first to measure reaction times for a numerical comparison task. Here, participants were presented with pairs of Indo-Arabic numerals ranging from 1 to 9 and were asked to identify the largest one by pressing a switch. The most interesting finding here is that both accuracy and reaction time varied in relation to the numerical distance between the displayed digits: the more numerical distance there was, the less time it took for subjects to tell which was the largest number.

For example, participants took more time to determine that 5 is larger than 4 than to say that 8 is larger than 1. In other words, when comparing numbers with respect to their size, performance increased with increasing numerical distance between two stimuli. This is known as the *distance effect*. Another important finding in this study was that reaction time was also influenced by the absolute magnitude represented by the integers: the smaller the numbers, the less distance was required to tell them apart. For example, response time was faster when comparing 2 vs 3 than when comparing 8 vs 9, despite the same numerical distance between the stimuli in both cases. This is known as the *size effect*. Together, distance and size effects are considered the behavioral signature of the system underlying numerical abilities beyond the subitizing range.

Such effects are not limited to Indo-Arabic numerals: distance and size effects also affect performance in comparison tasks involving number words.[56] For example, in Dehaene's (1996) study, subjects performed a comparison to standard task, which requires them to determine whether stimuli are smaller or larger than a standard (in this study, as is often the case, the standard was 5). To determine whether symbolic format has an influence on behavior, subjects had to complete the task using both Indo-Arabic numerals and number words for the numbers 1, 4, 6, and 9. Response times and

---

[56] Dehaene 1996; Ischebeck 2003; Lukas et al. 2014; though see Cohen Kadosh 2009 on how language-specific effects can modify the distance effect.

accuracy showed identical distance effects in both formats.

These performance effects are also visible in many other numerical tasks and formats. For example, subjects are quicker to reject false answers to arithmetic problems when the false answers are further apart from the true answer. Thus, when asked "7 + 5 = ?", subjects are quicker to reject 19 than to reject 13 (Ashcraft 1992).

Of course, most of the numerical stimuli we are exposed to in our daily lives go well beyond the numbers 1 to 9, so if distance and size effects are supposed to underlie numerical abilities beyond the subitizing range, there should be evidence of these effects in a much larger numerical domain. There is ample confirmation that this is the case. Distance effects have been found in numerical tasks using two-digit numerals,[57] fractions (Ischebeck et al. 2009; Jacob & Nieder 2009) as well as number words for two-digit numbers (Macizo & Herrera 2010).

2.2.2 Perceptual effects?

This being said, it should be noted that there is still some controversy about the influence of symbol organization, shape, and size for number words and unit-decade compatibility effects on performance in tasks involving complex numerical symbols. For example, while Dehaene and colleagues found "little or no discontinuity at decade boundaries" (Dehaene et al. 1992) for distance effects on two-digit comparison tasks, as Nuerk and colleagues (2015) report, processing is slower when comparing two-digit numerals in which symbols for units and decades are inversely related (e.g. comparing 47 and 62, where the decade relation, 4 < 6, is opposite that of the units, where 7 is

[57] Hinrichs et al. 1981; Dehaene 1989; Dehaene et al. 1990; for an excellent review see Nuerk et al. 2015.

larger than 2).

Early studies of two-digit numerals had minimized the importance of decade effects and concluded that two-digit numerals are processed by mapping the represented number onto an analog representation of discrete quantity (Hinrichs et al. 1981; Dehaene 1996). For example, Dehaene concludes that "our brain apprehends a two-digit numeral as a whole, and transforms it mentally into an internal quantity or magnitude." (Dehaene 2011, 64) In direct opposition to this interpretation, Nuerk and colleagues (2015; see also Nuerk et al. 2011) found many effects similar to decade-compatibility effects that they interpret as evidence for separate processing for decades and units. Such divorced processing suggests that we cannot simply extend findings from single-digit tasks to two-digit tasks and that we may have to amend the hypothesis according to which a single holistic analog representation underlies behavior in these tasks.

However, despite these notational effects, distance effects for multiple-digit numerals are still best explained as a combination of the combined distance effects of decades and units (Nuerk et al. 2015). Further, considering that distance effects were observed in congenially blind people performing a variety of numerical tasks (comparison, comparison to standard, parity judgment) using auditory symbolic stimuli for numbers 1-9 (Szucs & Csépe 2005) and for two-digit numbers (Castronovo & Seron 2007; Castronovo & Crollen 2011), the influence of format-specific visual effects on performance should not be overstated.[58]

[58] See Nuerk et al. 2011 for more a detailed presentation of various multi-digit presentation effects. Note that Nuerk et al. (2015) found that language influences the effect of decade-unit compatibility, with overall distance still being the most important predictor of behavior for German speakers, but not English speakers. According to Prior et al. (2015), "the lexical representation in a language influences magnitude comparison even when numbers are presented in a non-linguistic format." (85), though these effects may be due to phonological and graphical similarities of number words in certain languages (Lukas et al. 2014). Authors like Cohen Kadosh & Walsh (2009) and Campbell (2015) take such notation-specific effects as a sign that our representation of number is not abstract, as discussed below (Section 3.2) in

While overall distance does not account for all performance variation, it still accounts for much of it, which indicates that these findings do not undermine the existence of a format-independent representation of numerical magnitude displaying distance effects. As for number words, Macizo & Herrera (2010) found the same decade-unit compatibility effects in two-digit numbers presented in number word format. Interestingly, the same authors found facilitation effects for color naming tasks when objects depicted in stimuli shared phonological properties of the colors named (Macizo & Herrera 2014). Cohen Kadosh (2008) found that distance effects were different for number words presented in English and Hebrew, which Lukas and colleagues (2014) then showed could be explained by potential priming effects of graphical properties of number words: because some words share letters or morphemes, they are easier and faster to process than those that do not. Thus while these linguistic factors do affect performance, they do not negate the presence of a modality independent representation of discrete quantity with distance and size effects. This seems to support the existence of a modality-independent, shared internal representation for all external representations of numbers.

2.2.3 Distance and size: What do these effects mean?

Assuming we accept the data indicating that performance in numerical tasks is affected by numerical distance and size, the question is: why would there be *any* differences in reaction times and accuracy for numerical comparison tasks at all? Such performance effects in tasks involving symbolic stimuli is unexpected, given that symbols – especially numerical symbols – are associated with distinct representations that are easily distinguished from each other. For example, both 5 and 9 are larger than 4. On

---

relation to Dehaene's triple-code model.

the face of it, there doesn't appear to be any reason why it would require more time to determine this relation for 5 than for 9. And yet, controls requiring participants to press buttons to identify numerals shows that stimulus shape and identification cannot explain these effects, since the purely perceptual aspects of the numerals do not display any perceptual cues that the symbols co-vary with physical magnitudes (Buckley & Gilman 1974; Dehaene 1992).

This comes as no surprise: in terms of physical appearance, 5 is no more similar to 4 than 9 is, and none of these symbols has a shape that poses any particular processing hurdles for our visual system. Thus, there appears to be no evidence that the shape of the numeral influences response time, which suggests that distance and size effects here are not based on purely perceptual effects, but rather on the content associated with the numerical symbols. So while such performance effects may be surprising, given the precise semantic distinctions between numbers, the facts speak for themselves. These effects have been found over and over again in a wide variety of numerical tasks since Moyer and Landauer's seminal study.[59]

Given the limited influence of perceptual features, the most widespread explanation for distance and size effects in such tasks is that numerical symbols, regardless of their format, recruit the same analog representation of numerical magnitude, which shares many of the properties that characterize representations of physical magnitudes like size and luminosity.[60]

According to Moyer and Landauer, their results "strongly suggest that the process used in judgments of differences in magnitudes between numerals is the same as, or

---

[59] E.g. Forciciai & Park 2017; Lipton & Spelke 2003; Brysbaert, 1995; Dehaene & Akhavein 1995; Dehaene et al. 1990.

[60] E.g. Libertus et al. 2007; Moyer & Landauer 1967; Buckley & Gilman 1974; Dehaene 1992, 1997/2011; Gebuis et al. 2016; Lourenko 2015.

analogous to, the process of judgments of inequality for physical continua" (Moyer & Landauer 1967, 1520). The idea is that since performance effects on reaction time in the numerical tasks discussed above are not related to stimulus shape or other perceptual attributes, participants must be mobilizing an analog representation of numerical quantity when comparing symbol meanings. As Dehaene puts it,

> As far as physical appearance is concerned, digits 4 and 5 are no more similar than digits 1 and 5. Hence, the difficulty in deciding whether 4 is smaller or larger than 5 has nothing to do with a putative difficulty in recognizing the shapes of digits. Obviously, the brain does not stop at recognizing digit shapes. It rapidly recognizes that at the level of their quantitative meaning, digit 4 is indeed closer to 5 than 1 is. An analogical representation of the quantitative properties of Arabic numerals, which preserves the proximity relations between them, is hidden somewhere in our cerebral sulci and gyri. Whenever we see a digit, its quantitative representation is immediately retrieved, and leads to greater confusion over nearby numbers. (Dehaene 1997/2011, 63)

Since the same effects found in numerical comparison tasks are also present when comparing physical magnitudes, we can find out more about the representational system recruited in numerical tasks by looking at other systems that share its performance signature (i.e. distance and size effects).

2.2.4 The Weber ratio: a representational system with limited accuracy

It has been known for a long time that stimulus discrimination of many continuous magnitudes like weight, duration, length, loudness, total surface area, brightness, and pitch is impossible when stimulus variation remains below certain sensory thresholds. The first person to systematically study sensory thresholds was Ernst Weber, whose pioneering work in the 1830s led him to hypothesize that the change in intensity required to allow perceptual discrimination is proportional to the stimulus magnitude. Tough Weber's finding was based on his work on the ability to tell two weights apart, this finding was later extended and formalized by Gustav Fechner into a psychophysical law which he called Weber's law,[61] according to which the minimum amount of variation required to produce a just noticeable difference between two perceptual stimuli can be expressed by a constant ratio.[62] Weber ratios are used to describe the discriminability relations between stimuli in many sensory modalities, representing a variety of physical magnitudes.

It is easy to see that distance and size effects are corollaries of Weber's law. Since the variation in intensity required to produce a just noticeable difference is proportional to stimulus intensity, smaller intensities require less variation to allow discrimination. This means that the same perceptual variation will have more chances of being detected if applied to smaller stimulus intensities – the size effect. On the other hand, the larger the stimulus intensity, the bigger the difference must be in order to tell them apart, given that sensory thresholds are proportional to stimulus magnitude. This is the distance effect. Both these effects mean that, in many perceptual discrimination tasks, accuracy decreases in proportion to stimulus intensity, a property known as *scalar variability* (e.g., Cordes et al. 2001; van Oeffelen & Vos 1982; Mechner 1958).

Behavioral and neural data indicates that the underlying representation of stimulus magnitude in perceptual discrimination tasks is scaled non-linearly: as stimulus intensity increases, the representations become less distinct and there is increasing

---

[61] Since Fechner came up with this law, it is also commonly known as the 'Weber-Fechner law', though it was initially labelled Weber's law by Fechner, in honor of his mentor's work. Fechner went on to develop the field of psychophysics, which offers descriptions of the relationships between physical magnitudes and the perception of such physical magnitudes. He famously developed a law expressing the relation between stimulus intensity and sensation intensity, where sensation intensity is a logarithmic function of stimulus intensity, which is now known as Fechner's law. See Heidelberger 2004 for more on Fechner; see Ross & Murray 1996 for more on Weber.

[62] This ratio can be expressed as $\Delta I/I$, where ΔI is the difference in stimulus intensity required to produce a just noticeable difference,(e.g. a difference in loudness), while I is the intensity of the target stimulus (e.g. a certain decibel value).

overlap between these, thus making it more difficult to tell stimuli apart (Nieder & Miller 2003).[63]

2.2.5 Modeling distance and size effects

To account for distance and size effects in numerical abilities, two models have been particularly influential. One, the accumulator model (Meck & Church 1983; Gallistel & Gelman 1992), is an example of a mechanism that could count without words. Comparable to a car's odometer, the accumulator would be a cognitive structure made up of a mechanism that sends a constant signal, a gate that opens and closes for a fixed amount of time, and a recipient that holds all the signals that passed through the gate in a cumulative fashion. Thus, when presented with numerical stimuli, the accumulator would open the gate for each element of the perceived numerosity, allowing a fixed amount of signal to pass through for each, thus accumulating an analog representation of the perceived quantity. The final level would represent a quantity, though no counting routine was employed in the process. Here, numbers are represented by the continuous state of the accumulator. Each number would be an accumulated amount of signal having passed through the gate that corresponds with the perceived numerosity.

We can think of this like a small dam that lets out fixed quantities of water into a bucket: for each perceived element of a numerosity, the dam opens and closes, always the same amount of time, releasing water into the bucket. The total amount of water in the bucket

[63] Though there is disagreement on how to best formalize this compression of representations of stimulus magnitudes as they increase, this is a debate which, thankfully, lies well beyond the interests of the current review. What matters for us is that there is ample evidence that representations recruited in numerical tasks follow Weber's law, with representations of numerical magnitudes becoming increasingly noisy as magnitude increases, and thus that numerical representations are likely a form of magnitude representations with scalar variability.

represents the perceived quantity, though there was no counting involved. How such an accumulator would be implemented in a physical brain is still a completely open question (Dehaene 1997/2011, 19).[64] While this proposal has gained traction, "considerable evidence militates against the accumulator model" (Carey 2009, 132). In particular, an accumulator would predict that larger numerosities would take more time to be processed, since these would take more time to fill up. So far, there is no evidence that this is the case.

Another well-received metaphor for these effects is that our representation of quantities rests on an internal analog continuum on which numerosities and meanings of number symbols are mapped.[65] Each location on this so-called 'mental number line' corresponds to a specific discrete quantity. Importantly, the mental number line is not evenly split up. Rather, its structure is logarithmically compressed, so that the space between 1 and 2 is the same as that between 2 and 4, which means that larger numbers will be represented closer together, thus making our access to these fuzzier than for small numbers.

Support for the mental number line model comes from many sources.[66] One of these is what is called operational momentum, which refers to the fact that people tend to

---

[64] See Williamson et al. 2008, an article in which mice on speed (meth) seem to have faster internal clocks. The result is that, given a stimulus with numerosity of four, the rats' accelerated accumulator seems to allow more signal to come in, processing the input as having a higher numerosity of 6. See also Dehaene 2007.

[65] Galton (1880) was the first to describe human numerical representations as being ordered on a left-right line, while Restle (1970) was the first to explicitly tie in Moyer & Landauer's (1967) work to an analog line. See also Gallistel & Gelman 1992, 2000; Dehaene 2003; Cantlon et al. 2009; van Dijk et al. 2015)

[66] E.g. Priftis et al. 2006; Kaan 2005; Longo & Lourenko 2007; Cohen Kadosh et al. 2008, with a response from Verguts & Van Opstal 2008. See also Cordes & Meck 2013; Rugani et al. 2015; though see Núñez 2011 and Núñez et al. 2012 who have challenged the universality of this number line and its logarithmic format.

overestimate the results of addition while underestimating the results of subtraction, suggesting that these tasks involve moving along a mental number line in the direction of the relevant operation (Knops et al. 2009, 2013). Perhaps the most striking evidence is what is known as the SNARC effect, or Spatial–Numerical Association of Response Codes.[67] The SNARC was accidentally discovered by Dehaene and colleagues in a study where participants were tasked with determining whether numbers were odd or even using a button in each hand. Irrespective of which button was in which hand, responses for smaller numbers were always faster using the left button, while responses for larger numbers were always faster for the right hand.[68]

Since the initial parity-related experiment that led to the discovery of the SNARC, it has become an extremely active area of research, given the opportunity it affords of studying the association of number and space. In many paradigms probing the SNARC, magnitude information is irrelevant to task performance (e.g. determining parity, physical size or color of a numerical symbol, whether a number word starts with a vowel or a consonant, or even the presence of a phoneme in a digit's name). This has often been interpreted as a sign that numerical symbols automatically link to numerical magnitude representations that are spatially arranged along a mental number line.[69] Among the many fascinating discoveries made in relation to the SNARC, one worth mentioning here is that the direction of the SNARC appears to be culturally-determined, to the extent that it follows the direction in which one learns to write,[70]

---

[67] The original discovery is in Dehaene et al. 1993. Since then, the SNARC has been extensively studied. For reviews, see Wood et al. 2008; Hubbard et al. 2005, 2009; See also Viarouge et al. 2014.

[68] Note that the SNARC effect is related to the size of numbers used in the experiment. Thus if an experiment is only using numbers from 0 to 5, five will be processed faster when using the right hand.

[69] However, Cohen Kadosh & Walsh 2009 question whether automatic tasks recruit abstract representations on a mental number line.

[70] Dehaene 2011, 71. For more direct proof, see Ito & Hatta 2004; Zebian 2005; Shaki et al. 2012, though see Pitt & Casasanto 2014 who claim that SNARC direction can be modulated with finger counting. Göbel et al. 2011 has a useful review of the influence of culture and finger counting on the mental

though task demands may also influence this direction (Göbel et al. 2011). Because the SNARC provides direct evidence of an association between numerical quantities and spatial configuration, it is often interpreted as reifying the metaphor of the mental number line.

2.2.6 Non-symbolic tasks: an amodal representation of numerosity

I just showed how we can explain the presence of distance and size effects in symbolic number comparison tasks by appealing to the fact that the underlying representation follows Weber's law, which describes the psychophysical relation between perceptual representations and the stimuli that elicit them. As for non-symbolic numerical stimuli, the weber ratio should describe the relation between two stimuli such that we are just able to tell the difference between them based solely on the discrete quantity of objects perceived within each stimulus.[71] For example, if it a subject is able to tell the difference between a group of 8 dots and a group of 12 dots, but not 8 vs 11, 8 vs 10, or 8 vs 9, then the ratio of 12:8 (3:2) predicts that, when presented with a stimulus made up of 30 dots, the same person will have to be exposed to a stimulus of at least 45 dots to tell the stimuli apart.

One important prediction that follows from postulating an amodal representation of numerical magnitude underlying symbolic numerical skills is that all numerical comparison tasks, both symbolic and non-symbolic, should rely on this representational system, regardless of how the stimuli are presented. After all, if the underlying representation of numerical quantity is the same for both presentation

number line.

[71] Xu & Spelke 2000; Feigenson 2007; Nieder & Miller 2003; Izard et al 2008; Feigenson et al. 2004; Piazza et al. 2010.

formats, then similar performance effects should affect both symbolic and non-symbolic tasks, since the task requires processing numerical information in both cases, and our representation of numerical magnitude follows Weber's law. Thus, even though we do not process symbols for numerical quantities in the same way that we process non-symbolic numerical stimuli (see Buckley & Gillian 1974), if their use in numerical tasks recruits the same underlying representation of numerical quantity, the same performance effects should be observed. A further prediction from the hypothesis according to which the same representational system underlies performance in both symbolic and non-symbolic numerical tasks is that this representation is abstract, in the sense of being amodal: given that numbers are not modality-specific, unlike colors, say, the system responsible for numerical behavior should also be amodal.

As I discuss in this section, the general consensus is that both these predictions have been confirmed, over and over again, in many numerical formats and sensory modalities. Indeed, following Moyer and Landauer's pioneering work and the many studies replicating[72] their results, dozens of variations of their design have been used to probe the representational system underlying our numerical abilities, amassing more and more proof that numerical cognition is based on an analog representation of numerical magnitude with scalar variability. I limit my discussion to showing that distance and size effects have been found in visual and auditory modalities involving non-symbolic numerical tasks.

In the visual domain, many experiments have detected the presence of distance and size effects when subjects are asked to perform numerical tasks on dot arrays. For example, Buckley & Gillman (1974) first replicated Moyer and Landauer's findings for non-symbolic stimuli by exposing participants to stimuli that were composed of either Indo-Arabic numerals or dot arrays, organized in either familiar, regular, or

[72] See Buckley & Gillman 1974 for early replication sources.

irregular patterns. Like Moyer and Landauer, they found that reaction time was inversely proportional to numerical distance in both modes of presentation. Over the years, these distance effects have been found in tasks involving stimuli containing larger collections of dots.[73]

The same effects are observable when participants are asked to complete a precise number of actions without being allowed to count. Whalen et al. (1999) adapted methods used in research on numerical cognition in animals and asked subjects to press keys a specific number of times (between 7 and 25) or to identify the number of stimuli in random sequences of flashes (containing between 7 and 25 elements). By presenting stimuli too rapidly to allow counting, they found the same psychophysical signatures (i.e. distance and size effects) that had been found using the same methods in animals (Meck & Church 1983, discussed below).

Barth et al. (2003) further extended and solidified the hypothesis according to which the underlying representation of numerical quantity is amodal by comparing performance in crossmodal numerical tasks with intramodal numerical tasks. They asked subjects to determine whether the numerosities of sequences of flashes were the same as those of sequences of tones, with numerosities in each modality ranging from 10-30 items. Interestingly, performance was not significantly affected in such crossmodal numerical comparison tasks when compared to intramodal versions, suggesting that tasks in both modalities recruit the same amodal representation of numerical quantity. Barth and colleagues also asked subjects to determine whether numerosities presented in temporal format (i.e. sequences of flashes) were the same as

---

[73] Up to 30 dots: Libertus et al. 2007; up to 32 dots: Fornaciai et al. 2016; 10-50 dots: Barth et al. 2003; 100-400 dots: Fornaciai & Park 2017. Interestingly, Fornaciai and Park (2017) found that when dot arrays contain so many items that they become too dense to allow for individuation of each stimulus, Weber's law no longer applies. In this high-numerosity range (e.g., 400 tightly packed dots), they claim that "texture-density mechanisms might drive numerosity perception when the items become too cluttered to be individually recognized" (Fornaciai & Park 2017, 2).

numerosities presented simultaneously in a spatial format (i.e. dot arrays). No performance deficits were found when comparing crossformat to intraformat performance, once again suggesting that a unique, amodal representation of numerical magnitude was recruited in these tasks, rather than modality- and format-specific representations.

To sum up, the fact that performance in symbolic and non-symbolic numerical tasks across sensory modalities displays the same size and distance effects observed in perceptual discrimination strongly suggests that symbolic and non-symbolic numerical stimuli recruit the same analog representation of numerical content. The representational system recruited in such numerical tasks shares the same characteristics of the systems involved in inequality judgments of continuous physical magnitudes like luminosity and weight – namely, it follows Weber's law. Taken together, then, studies showing the same performance signatures within and across modalities when adults are prevented from counting suggests that our numerical abilities are grounded in a non-linguistic, modality-independent analog representation of numerical magnitude with scalar variability. This is the Approximate Number System. The question is: how do we acquire this representational system? Are we born with it, or do we need to learn it from others? In the next two sections I review evidence for the existence of this representational system in human infants and animals, thereby showing that it is both innate and evolutionarily ancient.

2.3 The approximate number system in preverbal infants

2.3.1 Empirical support

We have just seen that adults rely on an abstract analog magnitude representation of numerical quantity when performing numerical tasks. Given that the representational system recruited in adults shares characteristics of perceptual magnitude representations, it is not unreasonable to expect to find evidence that this representational system, like the ones involved in representing other magnitudes, is innate. Studies of pre-verbal infants allow us to test this hypothesis since they can probe the nature of number representations before teaching and enculturation set in. As I show in this section, evidence from numerical studies in infants suggests that the ANS is indeed innate.

While animal numerical abilities have been investigated in many species, research into infant numerical skills has been comparably sparse, with only few studies recently probing infant performance in numerical tasks outside the subitizing range. Despite the scarcity of research, four main sources of evidence for numerical abilities in infants (habituation, VOE, cracker choice, and manual search) all support the presence of the ANS in infants (Feigenson et al. 2004; Izard et al. 2008; Carey 2009).

Fei Xu and Liz Spelke (2000) provided the first solid evidence that preverbal infants are equipped with analog representations of numerical quantities. They used a VOE habituation paradigm to determine whether infant's numerical discrimination abilities are limited by the magnitude ratio signatures found in adults and animals. In this study, 6-month-old infants that were habituated to 8-dot displays dishabituated to 16-dot displays, while those habituated to 16-dot displays dishabituated when shown 8-dot displays, suggesting that they could distinguish between stimuli based on their numerosity outside the subitizing range. Note that infants habituated to 8 dots did not dishabituate to 12-dot displays, indicating that, at that age, the ANS' discriminability ratio is set somewhere between 1:2 and 2:3.

These results were then extended to larger numerosities, where 6-month-olds were successful at discriminating 16 vs 32 dots (Xu et al. 2003), but not 16 vs 24, again showing a discriminability ratio of at least 1:2, but less than 2:3. Sensitivity to numerosity change detection appears to develop with age: while 6-month-olds did not discriminate between 8-dot and 12-dot displays, 9- and 10-month-olds did, though they failed when numerosity ratio was 3:4, or 4:5 (Xu & Arriaga 2007).

Even newborns only a few hours old appear to have some form of numerical representations. In a variation of Starkey et al.'s (1990) crossmodal matching experiment (described above), Izard and colleagues (2009) familiarized newborns to auditory sequences made up of specific numbers of repetitions of syllables (e.g. "tutututututu" or "rarararara") and then exposed them to visual displays that either matched or did not match the numerosity of the auditory stimuli. Neonates looked longer at displays with matching numbers of objects when the ratio was 1:3 (i.e. 4 vs 12, or 6 vs 18 stimuli), but not 1:2, leading the authors to claim that their results "provide evidence for abstract numerical representations at the start of human life." (Izard et al. 2009, 10384). These cross-modality matching studies have been called the smoking gun of the ANS (Hyde & Mou 2017), given that non-numerical interpretations of the data in terms of perceptual confounds like total area and average size appear out of bounds to skeptics, since behavior here is based on representations integrating information from more than one modality.

Although the characteristic precision limitations of the ANS appear to be present very early in human development, the system undergoes a great deal of change over the lifespan. While newborns require a 1:3 ratio to detect a change in numerosity (Izard et al. 2009), by six months of age infants are capable of differentiating a 1:2 ratio change, and by nine months they dishabituate with a 2:3 ratio change (Xu & Spelke 2000; Lipton & Spelke 2004), as we just saw. Behavioral data from numerical comparison tasks involving non-symbolic stimuli indicate that precision in numerosity discrimination continues to improve from three years to thirty years of age (Halberda & Feigenson 2008; Halberda et al. 2012).[74] Studies of auditory numerical

---

[74] Interestingly, numerosity estimates are context sensitive: to calibrate our ability to estimate quantities, all it takes is exposure to a few numerical facts. For example, Dehaene (1997/2011) claims that if we are

discrimination skills confirm that ANS acuity improves over time, since 9-month-old succeed – that is, dishabituate – at discriminating 12 from 8 tones, but not 8 vs 10, while 6 month-olds failed at both (Lipton & Spelke 2003, 2004).

Similar ratio signatures as those found using visual stimuli were obtained involving auditory stimuli. Using the head turn procedure, an auditory variant of looking time procedures,[75] researchers habituated infants to sequences of tones and then exposed them to novel auditory numerosities. Controlling for non-numerical auditory cues such as tone duration, sequence duration, acoustic energy, and between-tone interval, they found that 6-month-old infants were able to discriminate 8 tones from 16, but not 8 vs 12 and 4 vs 6 tones (Lipton & Spelke 2003, 2004), displaying the presence of a weber fraction between 1:2 and 3:4 – the same found in studies involving visual stimuli for this age. This provides evidence that the same magnitude ratio limits performance in discrimination of both auditory and visual stimuli, and thus that the same representational system underlies performance in both modalities. Evidence gathered using similar methods with stimuli consisting of events (e.g. a puppet jumping) has revealed similar limitations, with 6-month-olds showing an ability to discriminate between 4- and 8-jump sequences, but not 4- from 6-jump sequences (Wood & Spelke 2005).

As I mentioned above, there is considerable evidence that the numerical representations

---

shown a display of dots and told (veridically) that it contains 200 dots, this will improve our ability to estimate quantities between 10 and 400 dots. Cappelletti et al. 2013 found that improvement in ANS acuity following training in one modality generalizes to other modalities, again pointing towards a single, modality-independent representation for discrete quantity.

[75] The alert reader might wonder how one does a looking-time study when the stimuli are a stream of tones. Basically, the sequences of tones are played from one of two speakers, each to a different side of the infant, and if the infants are interested in what is being played, they look to that speaker, continuing to attend to it even after the sequence is completed. Thus, looking times can reflect interest in this paradigm, and if infants notice a difference in the number of tones coming from a speaker, their attention is drawn to it. (Carey 2009, 126)

of the ANS are ordered in terms of numerical size on a mental number line. Confirmation that numerical representations are already in this format from a very early age comes from the fact that infants appear to be equipped with innate ordering skills. After habituation to sequences of images of dot arrays presented in increasing order, infants dishabituated to decreasing sequences (Brannon 2002), mirroring the data obtained using similar methods with rhesus macaques. Also mirroring macaque skills (Flombaum et al. 2005, discussed below) is evidence that some interpret as a sign that infants represent arithmetical operations on analog magnitudes. In a follow up study to Wynn's (1992a) probe into the arithmetical abilities of infants, McCrink & Wynn (2004) exposed infants to collections of objects being placed behind a screen (in addition tasks) and being removed from behind a screen (in subtraction tasks) – but this time, using numerosities outside the subitizing range (e.g. 5 + 5 = ? and 10 – 5 = ?). The key variable here was infants' looking time towards the screen when it was lowered to reveal possible or impossible outcomes. Infants look longer at impossible outcomes than possible outcomes, which can be interpreted as a sign that they can represent operations on numerosities (Carey 2009), even outside the subitizing range.

2.3.2 Continuity in development

This brief review of infant numerical abilities shows that preverbal infants react to numerosity variations outside the subitizing range, and that their discrimination abilities display the same limitations, regardless of modality of presentation. This provides considerable support for the existence of an innate amodal representational system of numerical magnitude with scalar variability, the ANS, suggesting that sensitivity to quantity is part of the human conceptual repertoire right from the start.

Given that there is evidence that the same system underlies numerical abilities in infants and adults, the question may be asked of whether or not this is the same system that is recruited in precise arithmetical skills. I have discussed evidence showing that symbolic and non-symbolic stimuli recruit the same system in adults WHERE. This, combined with the infant data I just summarized, builds a strong case that the numerical magnitude system we are born with is the one that is recruited later in life when we learn formal arithmetical skills. There is indeed considerable evidence that this is the case.

While the details of which brain regions are involved in which systems lie at a different, implementation level from the one I am interested in, which is more at the cognitive level, it is worth mentioning that a wealth of neuroimaging research "consistently point to regions in the mid intraparietal sulcus as the source of approximate number representations in both adults and infants" (Piazza 2010). This part of the brain displays increased activation when subjects perform tasks such as comparison of numerical magnitude and passive exposure to numerical stimuli in both symbolic and non-symbolic format (Piazza et al.2007; Piazza 2010). There is also considerable evidence that the ANS underlies more advanced arithmetical and mathematical skills. For example, Amalric & Dehaene (2016) found increased activation of the same language-independent intraparietal areas recruited when subjects are asked to complete simple numerosity-detecting tasks in expert mathematicians exposed to sentences from many fields of mathematics (geometry, analysis, topology, algebra, etc.), suggesting that the ANS serves as a building block for more advanced numerical skills. There is also evidence that individual differences in ANS acuity reflect mathematical skills, since performing well in estimating numerosities is a predictor of performance in mathematics classes, and that training in non-symbolic approximate number tasks like adding and comparing groups of dots improves performance in exact arithmetic.[76]

---

[76] For a review, see De Smedt et al. 2013. For more on the role of the ANS in predicting mathematical performance, see Hyde et al. 2014; Park & Brannon 2013; Halberda et al. 2008; Bugden & Ansari 2011, 2016. For a dissenting view, see Leibovich et al 20XX, discussed below IN SECtION BLAH, as well as

In short, there is converging evidence from behavioral and neurological studies that the ANS is involved in the development of advanced numerical skills, and that it is active from a very young age. Having briefly mentioned evidence for ontogenetic continuity of the role of the ANS in human numerical skills, I now review evidence for phylogenetic continuity by looking at a few studies showing the existence of analogue representation of discrete quantity in the animal kingdom.

2.4 Numerical skills in animals?

Studying animal cognition can yield important insight into the origins of our conceptual repertoire. For example, animal cognition research can help determine which cognitive abilities are evolutionarily shared and which are the products of culture. When people talk about possible arithmetical abilities in animals, the infamous case of clever Hans, a horse who was mistakenly thought to express correct addition results by tapping his hoof the appropriate number of times, never fails to come up. It serves as a cautionary tale in the literature on animal numeracy, warning researchers not to anthropomorphize the behavior of animals despite apparent homologies with human behavior, and to make sure that competing interpretations must be ruled out in order to allow us to attribute numerical abilities to animals. Despite this historical blunder, research into possible animal numerical skills has flourished, accumulating data on a wide variety of cognitive tasks in an even wider variety of organisms.

In the domain of numerical cognition, the number of studies that have found behavior indicating that animal are equipped with innate systems that allow them to react to variations in discrete quantities in their environment has grown to the point where

---

the commentary by Inglis et al. 2017.

numerical cognition is now one of the most studied higher cognitive functions in animals (De Cruz 2016:1).[77] This should not come as a surprise, since representing quantities of objects in their environment helps animals survive, say, by allowing animals to keep track of the number of predators (or prey) they are running from (or to). Recently, techniques used in the study of infant cognition (e.g. looking time procedure, habituation, VOE, and search tasks) have been applied to the study of animal cognition and this has yielded strong evidence that animals have rudimentary representations of quantity that display the same limitations as the systems found in human adults and infants. In this section I review some representative experiments, in in order to determine to which extent looking into our species' phylogenetic history can shed some light on the origins of our numerical abilities. To anticipate: I will show that the ANS is evolutionarily ancient.

Given that rodents are by far the most experimentally studied animals, it is not surprising to learn that one of the most important findings in the study of numerical abilities in animals came from observing rat behavior. While Porter (1904) and Koehler (1951) both conducted experiments geared at testing the numerical abilities of birds in relatively controlled environments, Mechner's (1958) experiments with rats were the first to probe animal numerical abilities in a proper laboratory setting. Here, rats were starved and put into confined spaces containing two levers. While lever B could release food when pressed, it would only release its treasure if lever A had been pressed a specific number of times. After trial and error learning, rat behavior eventually settled on pressing on lever A the required number of times before pressing on lever B, regardless of whether the number of required presses on lever A was 4, 8, 12, or 16. However, it is important to note that rats do not always settle on pressing the lever

[77] As testament to the incredible body of data amassed surrounding numerical abilities in animals, Agrillo (2015) catalogs a whopping 19 species of fish whose ability to discriminate discrete quantity has been experimentally studied.

exactly the right number of times. Rather, the accuracy of their pressing behavior diminishes as the cardinality of the number of required presses increases. In other words, rat behavior suggests they can respond to numerical information, but only approximately, since the range of number of presses associated with a target number – i.e. the incorrect number of presses for a target number – expands as the target number increases. For example, if the target number is four, rats can press the lever between 3 and 7 times, but if the target number is 16, then can press it anywhere from 12 to 24 times – clear signs of scalar variability, a telltale signature of the ANS.

Of course, when interpreting behavioral data from subjects that cannot report on the causes of their behavior, it is important to take into consideration competing explanations for behavior. In particular, when studying numerical abilities in animals and preverbal infants, it is crucial that experimenters try to control for the influence of non-numerical magnitude information on behavior.[78] For example, could we explain rat behavior here by appealing to the duration of the pressing behavior, assuming that the pressing rate is constant?

One way to control for such non-numerical magnitudes in this case is to vary the degree to which rats are deprived of water: the thirstier the rat, the faster it will press on the levers. Therefore, if the rat is responding to duration instead of numerosity, it would end up pressing more times on the lever if it was hungrier. Mechner & Guevrekian (1962) showed that such variations in starvation did affect the rate at which levers were pressed – but not the number of presses. This suggests that the rats are responding to numerical information in their environment. Rats often pressed the lever one more time than required. This can be explained by the fact that if the rat stopped pressing before the target value, no food would come out, so it was more prudent to add an extra

[78] Unfortunately, as I discuss below in section 2.5, there is a sense in which it is theoretically impossible to control for non-numerical confounds in numerical tasks, since number naturally co-varies with at least one of these, no matter how many controls are applied.

press and make sure they didn't under-count. However, if the rat presses too often, it wastes energy getting the food. Thus, the rats had good motivation to keep track of numerical information.

Meck and Church's (1983) study further tested whether rats were responding to temporal or numerical cues by training them to respond to sequences of two or eight tones in which number and duration were confounded, and subsequently varying the duration or the numerosity of these sequences. If rats were responding to temporal stimuli they would change the number of presses when duration was varied, while if they are responding to numerical stimuli they would change duration if the required number of presses were varied. Results of these experiments show that rats can generalize to both duration and discrete quantity, since they responded to variation in number when duration was kept constant, and they reacted to duration when number of tones was kept constant. Here too, rat behavior displayed scalar variability, confirming data obtained from Platt & Johnson (1971), who used slightly different methods and showed that rat responses displayed scalar variability for target numbers up to 24.

Of course, although it is difficult to ignore the evidence that rats represent numerosities, this ability does not guarantee mastery of an abstract concept of approximate quantity. To test this, Church & Meck (1984) used the same lever-pressing paradigm described above, but with a multimodal twist. Here, rats were first trained to press lever A if they heard two tones, and lever B if they heard four. After this auditory training, the same training method was applied, but this time with visual stimuli. After separate training in both modalities, the rats were exposed to novel stimuli in which both auditory and visual stimuli were mixed. The rats were immediately able to generalize their training to respond to stimuli with the same numerosity as the single-modality stimuli (e.g. two tones and two flashes). A few years later, Capaldi & Miller (1988) also obtained results that suggest that rats have amodal representations of discrete quantity, since they can keep track of total amounts of food received, even if that food takes different forms (e.g. 2 raisins, 2 corn). This data suggests that the system in charge of detecting numerosities can process information from many modalities to form one abstract, amodal representation of discrete quantity.

Cross-matching paradigms similar to the ones mentioned above in the infant section were applied to monkeys, obtaining similar results confirming the existence of a modality-independent representation of discrete quantity (Jordan et al. 2008). Here, monkeys were exposed to sequences of sounds on some trials, while on other trials they saw sequences of squares. These sequentially presented stimuli were followed by a screen in which two stimuli were displayed side by side, one containing a number of objects that matched the number of sequentially presented stimuli, another with a different number of objects. Monkeys had to press on the matching number across presentation format and modality to get a treat. Their performance was above chance, and they could generalize it to novel numbers of stimuli. Since these animals can match auditory numerosities to visual ones, it strongly hints to an amodal representation of numerosity with the same precision limitations found in humans.

In another seminal study, Brannon & Terrace (1998) trained rhesus macaques to respond to visual stimuli based on the number of items (e.g. dots, animal shapes, geometrical shapes) on a display. Researchers varied the size, shape, and color of the stimuli in order to control for non-numerical confounds (see figure 2.1). In the first phase, the monkeys were trained to order stimuli based on the number of objects displayed using images that contained at most four discriminable objects. Then, researchers tested monkey's ability to generalize their ordering behavior to displays containing up to eight stimuli. Success rates were well above chance, even when asked to order pairs of novel stimuli, suggesting that monkeys can represent numerosities and that they can learn ordinal rules on these numerosity representations. Brannon and Terrace interpret this behavior as a sign that monkeys "represent the numerosities 1 to

9 on an ordinal scale" (Brannon & Terrace 1998, 746). In other words, such data suggest that monkey's numerical representations are ordered on a mental number line, and thus that such animals have representations of relations between numerosities.[79]

Importantly, monkey performance in these tasks also displayed scalar variability, with accuracy decreasing in relation to increase in numerosity. Similar distance effects were replicated in later experiments involving displays containing up to 30 distinct stimuli (Cantlon & Brannon 2006). According to Brannon and Terrace, the fact that accuracy of performance in these tasks is a function of numerical size – a sign of scalar variability – means that the system underlying monkey behavior shows the same distance effects as those found in human adults and infants. Like most, they interpret the presence of such effects as a sign that numerosities are represented in an analog manner.

Despite the fact that, for the most part, I will not focus neuroimaging data, it should be noted that further evidence for the existence of an analog representation of numerosity in animals comes from the recent discovery of neurons tuned to individual numerosities in parietal lobes of monkey brains that are homologous to the horizontal Intra Parietal Sulcus (hIPS), which is largely considered to be the place where human brains process numerosity (Nieder 2005; Nieder & Dehaene 2009). This work shows that there are areas of animal brains that encode numerosity. Another interesting consequence to the discovery of single neurons tuned to specific numerosities is that we could possibly explain part of the increasing fuzziness of numerosity representations in terms of increasingly wide tuning curves for individual neurons. Given that tuning curves for numerosity-tuned neurons appear to widen with numerosity, we could explain part of the increasing imprecision of ANS representations by the fact that larger numerosities invoke wider activation patterns for neurons, which means that our representation of

[79] Brannon and Terrace point out that the mechanics underlying the monkey's ordering abilities is not yet known, but that one-to-one correspondence or counting could be candidates, depending on the 'operational' definition of counting used.

larger numerosities will involve a more confused activation pattern composed of the activation of many overlapping numerosity-tuned neurons.

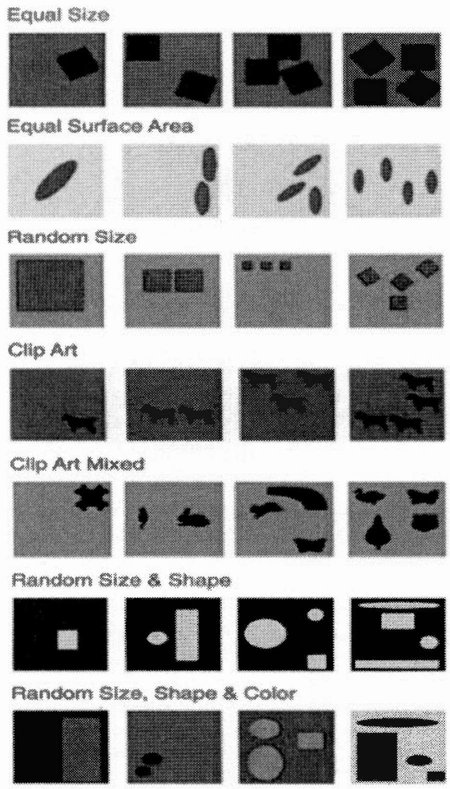

Figure 2.1 Variations in stimulus presentation from Brannon and Terrace (1998)

Finally, as was the case for infants, there is also evidence that many interpret as signs that animals can react to arithmetical operations on analog magnitudes. For example, when researchers exposed Rhesus macaques to four objects being placed behind a screen, followed by another four objects, looking time was longer if the screen was removed to reveal four objects than when the lowered screen revealed the expected eight (Flombaum et al. 2005). However, as discussed above, similar results obtained in infants can be reinterpreted as being caused not by variations in numerosity, but by changes in non-numerical magnitudes such as total surface area. For those weary of repeating the errors of the past and falling prey to what could be described as the Clever Hans effect, the same skepticism towards the methods used in infant studies can easily be applied to animal studies. To determine to which extent skepticism towards the methods used in studying numerical abilities in infants, animals, and adults can undermine evidence for the existence of an approximate number system, the next section summarizes concerns raised by those who doubt that it is possible to eliminate non-numerical explanations of behavior in such studies.

2.5 Approximate representations of numbers and analog magnitudes

In this chapter and chapter 1, I summarized data supporting the existence of two separate cognitive systems that underlie numerical abilities in human adults, infants, as well as in animals. Even if there are many outstanding issues concerning the systems involved in numerical cognition – e.g. how many systems there are, the role of attention, as discussed above, and which system plays what part in the development of mathematically-viable number concepts, as discussed in chapters 3 and 4 – the received view in the study of numerical cognition relies on both the OFS and the ANS to explain the development of number concepts. And yet, despite major progress in this field, a voice of dissent has been growing.

Given that numbers are generally considered to be as abstract and amodal as it gets, it is not always easy to isolate numerosity, described as the number of *perceived* objects, as the only potential cause of behavior, and it is usually not difficult to find a non-numerical explanation of behavior in many experiments. Such methodological skepticism has been around for a while,[80] as hinted at earlier in the discussion surrounding Karen Wynn's famous VOE experiment and its ability to establish numerical content in the ANS (section 1.4).

However, despite the methodological headaches that go hand in hand with numerical cognition studies, the consensus has generally been that researchers manage to find ways of controlling for non-numerical confounds in their experiments, and that results obtained from those who fail to do so will eventually fail to be replicated. So while there has been some skepticism about numerical interpretations of behavioral and neuroimaging data for a long time, it has not prevented ANS-based approaches to numerical cognition from becoming the most widespread interpretation of the data (Gebuis et al. 2016; Leibovich et al.2017).

Despite the widespread acceptance of the ANS as a legitimate explanandum of numerical behavior, a handful of authors have recently taken this skepticism to a higher level, questioning to which degree it is possible to create experimental conditions that can only be interpreted as evidence of the presence of innate cognitive systems with numerical content.[81] For example, in a statement that sums up the concerns flagged by this recent wave of ANS skepticism, Leibovich and colleagues have argued that

> the natural correlation between numerosities and continuous magnitudes makes it nearly impossible to study non-symbolic numerosity processing in isolation from continuous magnitudes, and therefore, the results of behavioral and imaging studies with infants, adults, and animals can be explained, at least in part, by relying on continuous magnitudes. (Leibovich et al. 2017, 1)

Typically, an ANS-skeptical argument goes like this: since it is impossible to control

---

[80] E.g. Simon 1997, Feigenson et al. 2002 Clearfield & Mix 1999; Uller et al. 1999.

[81] E.g. Leibovich et al. 2017; Gebuis et al. 2016; Lourenko 2015; Gebuis & Reynvoet 2012a; 2012b; Gebuis et al. 2014; Leibovich & Ansari 2016; Leibovich & Henik 2013; Soltész & Szücs 2014. For an early review, see Mix & Sandhofer 2007. See also Rips et al. 2008a.

for non-numerical cues in explanations of behavior in non-symbolic numerical tasks, the evidence for the ANS is dubious, and should be abandoned in favor of a more general Analog Magnitude System (AMS hereafter) whose domain is not specific to numerosity.

Thus, while there is no doubt that adult humans can produce estimates of the number of items they are paying attention to, there is controversy around the identity of the cognitive system responsible for this ability. Much of this controversy concerns whether the data surveyed in the previous section are best explained by appealing to a system specifically tuned to detecting quantities of discrete items (the ANS), or whether it is more prudent to appeal to a general sense of magnitude that is capable of responding to numerosity variations due to the fact that number co-varies with other magnitudes, in which case the system would be a more general analog magnitude system (the AMS).

Considering that data supporting the existence of the ANS and the OFS in infants and animals are based on using non-symbolic stimuli, this ANS skepticism threatens to undermine the enormous body of evidence supporting the ontogenetic and phylogenetic origins of the approximate number sense. To determine to what extent this criticism affects the strength of the dual-systems approach I have described above, in this section I sketch some of the main arguments levied against ANS-based theories. To simplify the discussion, I focus for the most part on a recent critical review representative of this new wave of revisionist skepticism (Leibovich et al. 2017).[82] In this section I briefly review some of the reasons offered to deny that experiments studying our numerical abilities reveal the presence of a system dedicated to

[82] While the details of the positive proposal offered by this review differ from some of the other voices associated with this skeptical approach, the criticism of mainstream numerical cognition methods offered here is shared by virtually all of these revisionist accounts.

representing quantities of discrete objects.

Nonplussed by the colossal progress made in the field of numerical cognition under an ANS-based framework, Leibovich et al. (2017) advocate overthrowing the dominant ANS-based theory and replacing it by an Analog Magnitude System (AMS), where numerosities and other continuous magnitudes are processed holistically. Their argument mainly turns on the impossibility of controlling for every non-numerical magnitude in non-symbolic numerical tasks, because number necessarily co-varies with one of these. To support this claim, they divide the methods used to control for non-numerical magnitudes in non-symbolic numerical tasks, and show that none of them can eliminate non-numerical interpretations of the behavior.

The first method is to manipulate a single continuous magnitude (e.g. by keeping it constant) while varying numerosity throughout the experiment. This way, since only numerosity varies, behavioral change should be due to numerosity alone, rather than to the magnitude that was kept constant. The problem with this approach is that there is no way to manipulate one continuous magnitude without affecting others. For example, in experiments that change numerosities while controlling for reaction to total surface area of the display by keeping it constant, numerosity variation necessarily incurs average size variation in objects.

A second method is to vary many continuous magnitudes throughout the experiment, though for each trial only one magnitude is manipulated. The same problem applies here, since for each trial, participants could be responding to different non-numerical cues, so that their performance can be explained by a variety of non-numerical strategies throughout the experiment.

A third method is to create congruency conditions between numerical and non-numerical magnitudes, in a Stroop-like paradigm adapted to numerosity. The idea here is to see if manipulating numerosity and an associated continuous magnitude have cumulative effects on performance. If there are no behavioral differences between congruent and incongruent trials, then the manipulated magnitudes do not interact with each other. For example, when asking participants to compare numerosity or area of dot arrays, congruent displays would be those where both size of dots and their numerosity are larger in one display than another. For example, in a study using such congruency manipulations (Nys & Content 2012), congruency effects were more marked for a size comparison task than a numerosity comparison task, which was interpreted by the authors as indicating that numerosity affected size comparison more than vice versa, and thus that numerosity is more salient in such tasks. The problem here, according to Leibovich et al., is that such results are difficult to replicate, and similar methods often support contradictory conclusions. In this case, a number of studies found, on the contrary, that numerosity was less salient than size (Leibovich et al. 2017, 5), which skeptics take as a sign of the task-dependence of many results of numerical cognition, thus indicating their unreliability. In short, the claim is that it is 'virtually impossible' to control for non-numerical explanations of behavior, since there is always a continuous magnitude that co-varies with numerosity in non-symbolic numerical tasks (see figure 2.2).

While the methodological humility advocated by many of these ANS skeptics is certainly warranted, especially considering the unfortunately liberal use of numerical terminology that permeates the study of numerical cognition, its limitations become apparent when confronted with stronger evidence for numerosity. For example, as mentioned earlier, a particularly strong line of evidence for the existence of the ANS is the cross-modal matching studies involving infants (e.g. Starkey et al. 1990; Izard et al. 2009; Feigenson 2011; Jordan & Brannon 2006) and animals (Meck & Church 1983; Jordan et al. 2008). As mentioned above, in such cases, it is difficult to explain behavior without appealing to an abstract, amodal representation of discrete quantity, since any modality-specific confounds like average object size or total surface area could not transfer across modalities, suggesting that amodal numerosity underlies behavior in these tasks. This would seem to nullify the numerosity skeptic's weapon of choice, perceptual confounds.

According to Leibovich et al., "Such evidence, however, should be taken with a grain of salt" (Leibovich et al. 2017, 5), since these findings have been very difficult to replicate, with only 2 of 6 studies managing to find evidence of cross-modal matching in infants. The problem here is supposed to be that since the findings are hard to replicate, they are worthless. And yet, while the small number of studies that managed to find evidence of cross-modal matching does highlight the difficulty of testing for the presence of numerical representations in infants, simply negating the original finding because it has not been easy to replicate appears unjustified. The same authors similarly question the validity of these studies by appealing to perceptual limitations in infants. For example, they claim that in many cases (e.g. Izard et al. 2009, mentioned above), the fact that infants have poorly developed visual acuity means that "they are unlikely to be able to see objects that are placed relatively close to one another as being separate from one another, and they lack the ability to separate between object and background or between one object and another" (Leibovich et al. 2017, 6). If this is true, infants in cross-modal matching tasks could be reacting to MORE/LESS cross-modal matches: hearing more syllables, the infant expects to see more dots, and thus stares longer at the matching stimuli.

However, according to a number of commentators, this is not true. For example, Hyde & Mou comment that "The claim that infants cannot perceptually individuate objects until 5 months is simply false" (Hyde & Mou 2017, 26). On the contrary, data suggest that while infant vision is poorly developed, it does not prevent them from individuating objects: "studies investigating newborns' visual perception have demonstrated that they are able to represent individual objects, at the same age as in the numerosity study" (De Hevia et al. 2017, 21).

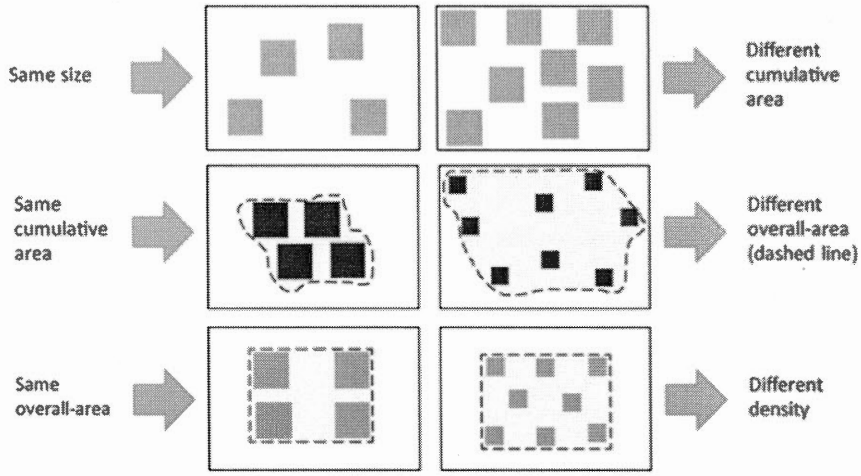

Figure 2.2 Correlation between numerosity and continuous magnitudes

There are many other problems facing such extreme skepticism with respect to what is now an incredibly diversified and well-documented body of research backing up the existence of the ANS. For example, given that the ANS skeptic's bread and butter is to appeal to the fact that number co-varies with continuous magnitudes to explain behavior, the same method might as well be used to deny the existence of any of these other magnitudes as well: why would numerosity be the odd man out, instead of convex hull? Worse, while it may appear plausible to deny that participants react to numerosity in many studies, one could be forgiven for being skeptical about the motley crew of magnitudes that has been recruited to replace it: how plausible is it that for the same numerical task (e.g. comparing dot arrays with respect to their numerosity), a variety of continuous magnitudes are recruited, depending on which one co-varies with number?

Also, given that which member of this motley crew of magnitude representations is recruited depends on the potential confounds with numerosity, wouldn't there be a system required to determine which magnitude should be recruited in each case to give the appearance of numerosity-based responses, and if so, wouldn't this system basically be the ANS? Last, consider that in a number of infant studies, infants are shown two stimuli and then allowed to reach for the one they want. Given that infants in such tasks are reaching using their hands to grab an individual object, it is difficult to understand how continuous magnitudes can underlie their behavior, given that hands do not grab convex hulls nor average luminosities, nor any other continuous magnitudes. Rather, given that hands are made to reach for things that can be grabbed in such manual search paradigms, there must be a discrete magnitude underlying behavior, and the only discrete magnitude that can account for behavior is numerosity. Thus, while it is important not to underestimate the potential influence of continuous magnitudes in numerical cognition studies, it is equally important not to overstate it, lest we only see the continuous forest and forget about the discrete trees.

A detailed discussion of the problems associated with such revisionist skepticism would take too much space here (for an overview, the commentaries on Leibovich et al.'s (2017) critical review provide a representative sample). Thankfully, there is reason to think that these skeptical concerns have minimal impact on the main topic of interest here: considering that this chapter is aimed at identifying and characterizing the representations that underlie numerical cognition, the question of whether the system that allows us to represent numerosity is dedicated to numerosity alone or to both numerosity and continuous magnitudes is of secondary importance. Despite the concerns raised by ANS skeptics, the behavior of the system in charge of processing numerosity – in particular, its precision limitations due to the fact that it follows Weber's law – is not put into question, regardless of whether it takes the form of an ANS or an AMS. Similarly, given that my interest here is to understand how our developed number concepts grow out of our innate cognitive machinery, the precise neural instantiation of the system is of secondary importance.

The important point is that, on its own, the system involved in numerosity processing cannot explain the emergence of mathematically-viable number concepts. The same can be said for the system responsible for numerical skills inside the subitizing range, covered in the previous chapter. The question that interests me is how these limited systems are involved in the formation of more precise representations of numbers. On the issue of how limited systems like the ANS/AMS and object files allow the construction of natural number concepts, the ANS and AMS approaches need not disagree. So while it may indeed turn out that numerosity is extracted from generalization on other magnitudes rather than being the output of an innate representational system that encodes numerosity, the important point is that there are systems that allow us to respond to numerosity, and that the limitations of these systems prevent them from giving us the whole story on where precise number concepts come from. Now that we have a good idea of what these systems look like, we can take a look at how they can be used to construct mathematically-viable number representations. This is the aim of the next two chapters.

2.6 The Gap Problem in numerical cognition

The empirical data reviewed here and in the previous chapter established that animals, infants, and adult humans share some cognitive systems that allow them to represent quantities of discrete objects in the environment. While there no conclusive evidence that the OFS has any representations with numerical content, it may still underlie behavior based on numerical properties of stimuli via one-to-one correspondence between individuated object files. Importantly, its range is severely limited. On the other hand, the ANS's range shows no such limitations, but the fact that its precision decreases in relation to increasing numerosity in conformity with Weber's law means that, on its own, it does not generate representations with sufficiently precise content to be described as numbers. Despite their limitations, however, the evidence considered from behavioral and developmental studies as well as neural imagery all display the signatures of these systems, which suggests that they continue to operate throughout our arithmetical lives and underlie our basic numerical abilities to count and perform basic arithmetical operations.

Considering the severe limitations of these systems compared to the precise character of the objects of mathematical practice, there is a lot of work to do before we understand how humans acquired precise representations of numbers. In particular, how these systems interact with each other, with other systems, and with the environment, both in ontogeny and phylogeny, remains unclear. While both the ANS and the OFS are involved in representing information about quantities of things in our environment, their obvious size and precision limitations suggest that a major change must occur in order to allow us to develop the fully-fledged number concepts used in arithmetic, whose precision and infinite extension cannot be accommodated by either innate system. Without such a major change, it is mysterious how we could even come to think about the number six, say, given that neither system can explicitly represent precise numbers larger than four. There is thus a significant gap between the content generated by the neural systems we are born with, and the content we use in precise arithmetical thinking.

This is what I will refer to as the *gap problem*: how do we bridge the gap between the limited numerical content produced by our evolutionarily ancient brains and the properly numerical content associated with numeration systems like Indo-Arabic or Roman numerals? There are many different accounts of how we managed to sharpen the approximate representations inherited from our evolutionary lineage. In the next two chapters, I discuss some of the most prominent accounts of how we bridge the gap between our limited innate cognitive machinery and the number concepts we use in mathematical practice.

CHAPTER III

DEHAENE'S NUMBER SENSE

3.1 Introduction

The empirical data reviewed in chapters one and two conclusively established that animals, infants, and adult humans share some cognitive systems that can represent collections of discrete objects in the environment. Considering the severe limitations of these systems and how this contrasts with the precise character of the objects of mathematical practice, there is a lot of explanatory work to do before we understand how humans acquired number concepts. In particular, how these systems interact with each other, with other systems, and with the environment, both in ontogeny and phylogeny, for them to be recruited in arithmetical tasks, remains unclear. There are two main questions concerning the relation between systems like the ANS and the OFS and our advanced numerical abilities: what sort of cognitive architecture allows the processing of numerical information, and what sort of process allows this architecture to develop from our evolutionarily-inherited systems? There are many competing accounts of how we managed to sharpen the approximate representations inherited from our evolutionary lineage. This being said, despite the tremendous variety of such accounts of how numbers could come from processes in our heads, the most prominent account of the development of numerical cognition must be Stanlislas Dehaene's number sense.

No single researcher is more associated with the field of numerical cognition than Dehaene, a mathematician-turned-neuroscientist. Prior to the first publication of his popular science book, *The Number Sense* (1997/2011), organized research into the cognitive roots of our mathematical skills had barely taken off, with only a few labs starting to exploit the then-new methods and technologies of various fields of cognitive science to probe the neural underpinnings of numerical abilities in humans and animals (Dehaene 1997/2011, ix-x). Since then, numerical cognition research has grown by leaps and bounds, to the point where it is now an important domain of research in its own right. Dehaene is easily one of the most important figures in this field. His name almost inevitably shows up in the reference list of any philosophical work dealing with numerical cognition, as well as most articles dealing with the relation between our brain and various aspects of arithmetical practices.

As a testament to Dehaene's lasting influence, consider the fact that his triple-code model of how our brain represents numbers, first explicitly formulated in 1992, is still considered "the currently most influential model in numerical cognition research" (Link et al. 2014, 1),[83] more than twenty years later – centuries in the fast-developing field of numerical cognition. In this chapter, I summarize the main lines of Dehaene's account of the historical and ontogenetic origins of numerical cognition. Along with Susan Carey's bootstrapping account, discussed in the next chapter, Dehaene's account is one of the most detailed explanations of how the systems discussed in the first

---

[83] Similarly, Siemann & Petermann (2018) describe the long lasting influence of Dehaene's Triple Code Model (which they abbreviate here as 'TCM') as follows:

> Since TCM was first outlined in the late 1990's, there has been major progress in terms of methodology (e.g. brain imaging), developmental studies (cross-sectional; longitudinal) as well as diagnostics. Therefore, some of the initial statements were elaborated on, others revised and some rejected. Nevertheless, the main scaffold of this multiroute model still shapes the majority of arithmetic theories, and findings in various fields of arithmetic research today are frequently integrated into the framework of TCM. (Siemann & Petermann 2018, 114)

For an recent fMRI analysis of the triple-code model, see Skagenholt et al. 2018.

chapters are modified in order to allow the development of representations with what could be described as proper numerical content (in opposition to the quantity representations of the unadulterated ANS). Thus, Dehaene is the ideal starting place for a discussion of how natural numbers could be the product of cognitive systems like the ANS and the OFS.

In the first part of this chapter, I summarize the main lines of Dehaene's model of the cognitive architecture responsible for our numerical abilities. I then turn to his account of the culturally-induced neuronal recycling that maps number words and numerals to representations of the ANS which he claims explains how we develop advanced numerical skills, before presenting Susan Carey's (2009) criticism of this mapping in the final sections of the chapter.

3.2 Modelling the cognitive architecture for number

3.2.1 Abstract numbers and numerical format

The question of how the brain comes to represent number poses a unique challenge for researchers, given the abstract properties of numbers and the material and fleeting aspect of brain activation that is often appealed to when attempting to ground abstract objects in material reality. As Dehaene put it, "Research in this area tends to cross the traditional boundaries of cognitive science" (1992, 2). Given that we can apply number concepts to any discrete object of thought, they look like the perfect candidates to support the existence of abstract, amodal representations. And yet, despite their abstract nature, numbers take a variety of forms in our lives: we read and write numerals and number words in a variety of formats, we hear names of numbers, and we see and hear non-symbolic numerical information (e.g. we see a number of spoons to set the table, or hear beeps before the start of a race). Numbers present themselves to us in many forms. This includes representations of numbers in linguistic format, as is spoken and written number words, in non-linguistic symbolic format, like the Indo-Arabic and Roman numerals, and non-symbolic format, as in the case of collections of dots or sounds. Despite this variety of modes of presentation, there is something common to the various ways in which numbers are represented in the practice of arithmetic – namely, they are representations of *numbers*.

A question that naturally surfaces when considering the variety of forms in which numerical information is presented to us is whether or not the various formats in which we represent numbers share neural resources. Of course, as Campbell (2015) remarks, the content of a representation need not be reflected in its cognitive implementation: "number and arithmetic are abstract when viewed as formal mathematical concepts, but does this extend to the cognitive domain?" (Campbell 2015, 140) The issue of how the brain encodes numerical stimuli has implications for the extent to which representations of numbers can be considered to be amodal: if the brain were to encode numbers via an abstract code shared by every one of its formats, then this could count as support for seeing numbers as abstract[84] entities. If different formats and modalities recruit entirely different parts of the brain, this would support seeing number as a patchwork of representational content.

If the brain has a single, amodal representation of numbers, this could explain why it is possible to assign a number to any discrete amount of things we can think about, whether it be sounds, images, touches, or any other modality. If, on the other hand, numbers are encoded in format-specific neural code, then we might expect to have

[84] While the notion of abstractness is important for this discussion, it is not necessary to get a detailed account of what it means here, given that our interests lie far from such metaphysical matters at this point. For present purposes, we may characterize the use of the word 'abstract' as applying to content/representations that is not domain-specific, that we can process identically independently of format or modality of presentation. See Campbell 2015 for an empirically-oriented discussion of abstractness in arithmetic. See also Cohen Kadosh & Walsh 2009.

more difficulty processing numerical information in some modalities when compared to others. This would be due to the fact that there would be more translation required to get from one format to the other, given that a hypothetical 'core' format of numerical representation would be more easy to access for some modalities than others.

Consider what happens when we are presented with a multi-modal numerical task – say, hearing someone ask us to calculate something, and writing down the answer to the verbal query on a piece of paper. For this to happen, the verbal command must be encoded in a specific format, and then we must either calculate the answer or recall it from memory (for simpler calculations) before accessing the code responsible for producing written symbols. If the neural code for number is amodal, then the calculation will take place in the shared neural code before being translated into the output modality. If, on the other hand, neural code is a patchwork, then there can be modality- and format-specific processing that does not require an amodal middleman for such translation. One question that relates to this is the extent to which numerical processing stages are interactive or additive: if they are interactive, then we can expect there to be a relation between the format in which a task is presented and the resources recruited by that task, and the total processing for the task will be a function of the level of interaction between the recruited formats or modalities. If they are additive, then every format and modality processes number independently of other encodings, and tasks that require multiple formats or modalities will recruit each in turn, with total processing for the task being the sum of each individual processing stage.

Three main models[85] have been proposed to describe the cognitive architecture underlying the processing of the various ways numerical information is presented to us, and the extent to which numerical processing relies on an abstract representation of number: Dehaene's (1992) triple-code model, McCloskey's (1992) abstract code

[85] See Cohen Kadosh & Walsh 2009 for references to less discussed models, including their own dual-code model.

model, and Campbell and colleagues' (2004, 1988) Encoding Complex Hypothesis. Each model differs from the others in terms of the extent to which the representations that underlie our responses to variations in presentational format and modality are abstract. To understand Dehaene's answer to how we build number representations out of the systems explored in chapters 1 and 2, it is important to familiarize ourselves with the main lines of his triple-code model and how it differs from the other two main models. Familiarizing ourselves with the triple-code model allows us to see how data at the implementation level constrains possible cognition-level descriptions of which architectures could implement numerical computations in human brains. Such models of cognitive architectures can then help identify what sort of learning processes allow the development of numerical cognition.

3.2.2 McCloskey's abstract code model

McCloskey and colleagues' abstract code model could perhaps be qualified as the most intuitive, given that the basic idea here reflects the commonly-held view that numbers are abstract entities, and that this means there should be an abstract representation of numbers somewhere in the brain. On this model, every numerical task taps into an abstract number representation via a modality-specific translation process. Numerical inputs from different modalities and formats are converted to an abstract semantic representation (McCloskey& Macaruso 1995) of numerical magnitude which then serves as the basic form of the output, be it in the production of Indo-Arabic, verbal, or non-symbolic numerical representations. Notation-specific comprehension and production modules mediate the input and output between the symbols and the amodal representation of number. On this account, calculations take place in the amodal representation, which means that any task involving more than one format involves 'asemantic transcoding', to use McCloskey's terms, between the two formats in the central abstract representation. Processes that require input from one format and output via another must then pass via this central abstract representation of number, with the output being the result of an additive process of perception, translation, central processing, and further translation to the output modality (See figure 3.1). Here, the fact that processing time can vary between modes of presentation of numerical information is interpreted as a sign that some translation and production processes take more time than others.

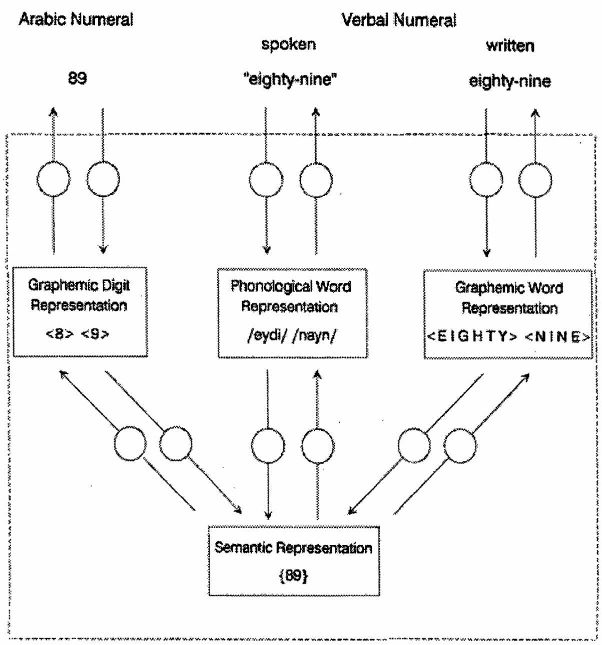

Figure 3.1 Abstract Code Model (From McCloskey & Macaruso 1995)

The main support here comes from lesion studies that show that some arithmetical processes can be impaired while others spared, which suggests that no format has privileged access to numerical content (McCloskey & Macaruso 1995). If we recruit different neural resources for different tasks and these interact in additive ways, this can explain why lesion studies selectively impair tasks, since a lesion can occur in a sub-component whose processing does not affect other tasks. For a long time now, lesion studies have shown that arithmetical abilities are doubly dissociated from each other, in that it is possible to be proficient in processing numerical information in one format (say, spoken number words) while being unable to process numerical information in another format (say, Indo-Arabic numerals).[86] There can also be intra-format disabilities. For example, some patients display production errors in parts of their mental lexicon for specific quantities (e.g. hundreds, tens, units, etc.), which suggests that the linguistic organization of number representations can be selectively impaired. Others suffer from problems with number syntax, as evidenced by certain aphasic patient's inability to properly code verbal numerals into their corresponding Indo-Arabic notation. Similar syntax- and lexicon-based errors are also observed in children learning number words (Dehaene 1992).

Further support comes from patients who are unable to produce the correct answer to arithmetical problems, but can nevertheless recognize the answer from a visually-presented list, and patients with the opposite ability (i.e. can produce an answer but cannot recognize it) (McCloskey & Caramazza 1985). This suggests a dissociation between number production and comprehension abilities. Also consistent with this model are data described in chapters 1 and 2 which find the presence of similar size and distance effects in all modes of presentation of numerical stimuli, as well as the fact that the Space-Number-Association-Response-Code (SNARC; see section 2.2.5) applies to verbal and non-verbal numerical stimuli, which suggests that these effects are due to the presence of the same underlying representation for all formats and modalities.

The main problem with this view is that there is evidence going against it. For example, consider the size congruity effect, according to which numerals whose physical size is congruent with their relative size in numerical magnitude comparison tasks will be

[86] See Dehaene 1992 for a review. See also part VI of Cohen Kadosh & Dowker 2015.

easier to process than when physical size and relative magnitude are incongruent.[87] Like the SNARC, the size congruity effect is generally taken as evidence that some aspects of numerical representations are processed automatically, since we cannot help but process numerical magnitude information despite the fact that it is irrelevant to the task.[88] However, Ito & Hatta (2003) found that such effects are absent from Kana verbal scripts for numbers, but not Kanji script, which is ideographic, not verbal. Cohen Kadosh & Walsh (2009) interpret this as evidence that format of presentation influences processing and automaticity, and thus that number is not encoded in an abstract, amodal format.[89] Further evidence against the abstract code model is the evidence favoring the encoding complex model, which I summarize in the next sub-section.

3.2.3 Campbell's Encoding Complex Hypothesis

In contrast with McCloskey's amodal representation of number, Campbell (2016) and colleagues (1988, 2004) have proposed the *encoding complex hypothesis*, according to which numerical stimuli activate a network composed of many different format- and modality-specific encodings. Here, there is no central representation common to all

---

[87] This effect is present in stroop-like conflict paradigms in which the physical size of numerals is varied as subjects perform tasks on these. For example, subjects can be asked to report on the comparative physical size of numerals and to ignore the magnitude they represent, or vice versa. Symbol size can be congruent with its relative size, as in (5 $_2$), neutral, or incongruent, as in ($_5$ 2).

[88] Adopting a definition of automaticity from Tselgov et al. 1996, Cohen Kadosh & Walsh (2009) describe a process as automatic "if it does not need monitoring to be executed" (317). See also Macleod & Dunbar 1989, who propose a continuum of automaticity.

[89] See Cohen Kadosh & Walsh 2009 for more detailed criticism based on automatic processing of numerals, as seen in size congruity paradigms. Myers & Szücz (2015) showed that reaction time for some calculation tasks varies depending on format of presentation. See also Campbell 2015 for evidence against interpreting differences in processing times between Indo-Arabic numerals and number words as being attributable purely to modality-specific translation differences, as McCloskey proposes.

formats of presentation of number. Rather, each format is represented with its own specific code and interacts with other codes directly (see Figure 3.2). Numerical representations here are not amodal, but multi-modal. Because there is no central representation of number on this model, each task will recruit relevant and irrelevant neuronal resources, given that any encoding is part of an intertwined complex of format-specific encodings that cannot be completely dissociated from one another.

Different presentation formats display differences in performances. For example, comparing the relative size of numbers presented in verbal format is slower than in symbolic format (Ischebeck 2003), while calculation with number words can be up to 30% slower than with Indo-Arabic numerals (Campbell 2015). Such effects of notational differences in simple arithmetic tasks (e.g. 4 + 5 vs "four plus five") are one reason behind Campbell's modality-dependent account of numerical cognition, although, as mentioned above, the abstract code model can account for such effects by appealing to format-specific differences in translation processes.

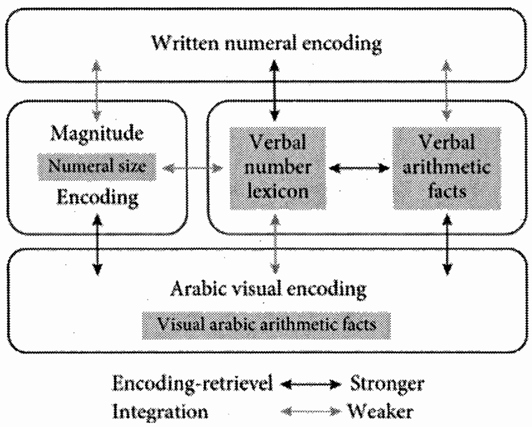

Figure 3.2 Encoding Complex Hypothesis (From Campbell 2016)[90]

[90] Note that in this schema, "Darker arrows correspond to greater encoding-retrieval integration, which means more automaticity and a stronger capacity to coordinate task-specific processing associated with each code." (Campbell 2016, 144)

Perhaps stronger support comes from the fact that numerical tasks often recruit more than one format automatically, despite the irrelevance of the information contained in the secondary format. Evidence for this comes from the above-mentioned cases of size congruity effect as well as apparent interaction-based errors in calculations. For example, it takes more time to say that 6 + 4 = 24 is false than 6 + 5 = 15, since the first result would be true if it were a multiplication task, which suggests that the multiplication result was recruited involuntarily. This automatic and involuntary retrieval of irrelevant information is interpreted as evidence that there is interaction between the different resources responsible for processing format-specific numerical content: "In the encoding-complex view, communication between representational systems often involves interactive, rather than strictly additive processes" (Campbell 2015, 144). This goes against the abstract-code model, which predicts additive interaction between different processes that operate independently from each other. Preaching for his choir, Campbell writes

> The evidence runs contrary to the view that arithmetic is essentially an abstract process that operates independently of encoding context or response output conditions. Instead, the evidence points to a cognitive architecture in which problem encoding and calculation processes are highly interactive and where linguistic codes provide an important, but not exclusive medium for arithmetic. (Campbell 2015, 140)

Unsurprisingly, however, this model does not fare as well with the main data supporting the abstract code model, including distance and size effects as well as the SNARC. Also unsurprisingly, there are ways of re-interpreting all this data that make room for the abstract code model.[91] This being said, both models have support going for them, and the matter of how the brain encodes numerical information is not settled (Myers & Szücz 2015), with evidence favoring and going against the amodal and the multi-modal option. The purpose here is simply to illustrate the commitments of each

[91] See the commentary section on Cohen Kadosh & Walsh 2009.

model and what sort of data can be used to support each, in order to highlight the sources of Dehaene's triple code model, to which I now turn my attention.

3.2.4 Dehaene's triple-code model

Building on elements of both the encoding complex hypothesis and the abstract code model, Dehaene's triple-code model includes both modality-specific processes and an abstract, amodal representation of numerical magnitude. This model relies on two main premises: 1) numbers can be represented in three separate neural codes, and 2) each numerical task has its own distinct input-output encoding. Dehaene's triple-code model, as its name suggests, divides the way in which the brain encodes number representations into three main categories: a visual code for Indo-Arabic symbols, an auditory verbal code for number words, and an analog magnitude code that contains what both Dehaene and McCloskey label a semantic representation of number, to describe the fact that this encoding processes information about the magnitude of a numerical representation. As in Campbell's model, each format here recruits its own specific code. However, as in McCloskey's model, there is a central, amodal representation that can be accessed by all modalities for tasks that do require understanding the magnitude associated with a symbol. Thus the visual code will process Indo-Arabic numerals (e.g. parity evaluation, symbolic calculation), the verbal code will process number words (e.g. arithmetical facts stored in verbal format, such as addition and multiplication tables), and the analog code will process magnitude information required for comparison and approximate calculation, which can take input from many modalities and formats (see Figure 3.3).

The reader familiar with chapter 2 will remember that the analog magnitude code is housed in the neural architecture of the ANS in the parietal lobes. Imaging studies are "consistent with the hypothesis that the hIPS codes the abstract quantity meaning of numbers rather than the numerical symbols themselves" (Dehaene et al. 2003, 492).

Surprisingly, although the ANS provides semantic representation of number, on Dehaene's model, it is not necessary to recruit this representation for all tasks. Rather, only those tasks that require such a semantic representation for their completion will involve the ANS. Purely syntactic tasks, such as calculation and retrieval of arithmetic facts, need not involve the ANS.

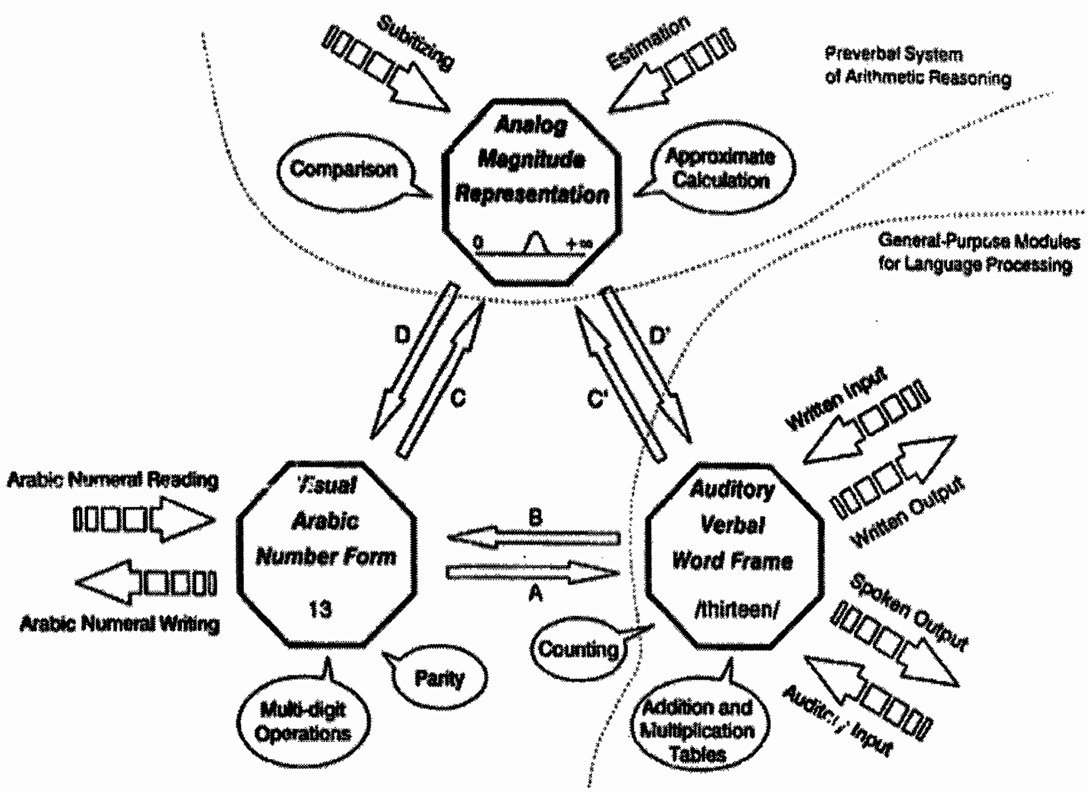

Figure 3.3 Dehaene's Triple-Code Model (From Dehaene 1992)[92]

One advantage of this model over its alternatives is that other models focus mainly on calculation tasks, leaving behind comparison, estimation, and quantification tasks,

---

[92] Note that Dehaene changed his mind, as mentioned in chapter 1, and that he no longer takes subitizing to be based on the semantic representation in the ANS.

which recruit different, evolutionarily ancient systems. According to Dehaene, this means that competing models are better suited for describing the syntax, but not the semantics of numerical cognition, given that he, like many (e.g. Menary 2015 and Dutilh Novaes 2013, discussed in chapter 6 and chapter 5, respectively) claims it is possible to carry on calculations without having an understanding of the meaning of the symbols we are manipulating.

On Dehaene's model, calculation recruits separate resources from arithmetical fact retrieval. This means that tasks involving both formats will require some form of format switching or re-coding. This should increase performance time, and does indeed appear to be the case in parity evaluation tasks when the numbers are presented in verbal format (Dehaene et al. 1993).

Out of the three representational systems responsible for the encoding in Dehaene's model, only one is language-based. It underlies numerical tasks that are presented in linguistic format, such as verbal counting and accessing arithmetical facts stored in verbal format in memory. According to Dehaene, the system responsible for behavior in such tasks is not specific to numbers. For example, while retrieving facts learned from memorized multiplication tables requires general memory skills involving verbal associations, verbal counting depends on the ability to learn lists of symbols by heart, which infants can also do with the alphabet, nursery rhymes, and other memorized sequences of symbols. We see then that for Dehaene, language plays a part in the development of more advanced number concepts, with the ANS and other such systems serving as the underlying representation of quantity.

According to Dehaene, a separate system is dedicated to completing tasks involving the symbolic representation of numbers, including Indo-Arabic numerals and other non-linguistic systems of numerical notation. This includes tasks like "multi-digit calculation or parity judgment which require the mastery of a dedicated positional notation system" (Dehaene 1992, 34).

Finally, for tasks like numerical comparison and approximation tasks, Dehaene describes the underlying representational system as being dedicated to number, in that it responds to information about numbers of objects in the world, and allows us to compare numerical magnitudes. These tasks can be performed by preverbal human infants and animals, which strongly counts against a linguistic origin to the associated abilities. As Dehaene sees it, the system that allows us to perform such tasks may embody "the main, and perhaps the only, 'semantic' representation of numbers." (Dehaene 1992, 35)

If the triple-code model is right, there is a sense in which there is no such thing as a single, unified, amodal representation that underlies our performance in all numerical abilities. Rather, different numerical tasks will recruit different neural assemblies:

> According to the proposed model, the ideal of a unique "number concept", which motivated Piaget's (1952) or Frege's (1950) reflections, must give way to a fractionated set of numerical abilities, among which faculties such as quantification, number transcoding, calculation or approximation may be isolated. (Dehaene 1992, 34)

According to Dehaene's triple-code model, our numerical abilities are spread into separate modules dedicated to processing specific format of numerical content. This model groups numerical tasks according to the format in which the tasks are taking place.

The fact that some forms of numerical representations are encoded by modality-specific systems means that some of our understanding of number is not purely abstract, since some tasks can only be completed by recruiting language and vision-based neural code, on this account. This being said, since this model also appeals to a modality-independent representation of numerical magnitude to ground these components, there is a sense in which some numerical content is modality-independent. Note that the presence of an amodal encoding in this model supports an additive interpretation of the interaction between its components, since once the stimulus is transcoded into the appropriate format, processing is format-specific and independent of other formats. This sets it apart from the encoding complex hypothesis.

This section has motivated Dehaene's triple code model by contrasting it with competing models and highlighting its similarities and differences with each. Whether or not Dehaene's triple-code model is right is an empirical matter that will hopefully be resolved in due time. For current purposes, what is important to get from this discussion is how data can constrain models of the architecture behind numerical cognition, and what sort of components play which part in the processing of numerical information. The triple-code model takes into consideration considerable data and builds with these a patchwork of neural systems that seem to underlie our ability to process numerical information in symbolic format.

While knowledge of what sort of system might be responsible for which arithmetical tasks is certainly an important piece of the puzzle of where our number concepts come from, there is an important piece missing from the puzzle here: how does our brain manage to link up the semantic representations of the ANS with other encodings for numbers? After all, even if the triple-code model is the right one, this fact alone does not tell us how the ANS or the OFS are modified or associated with non-semantic encodings of the triple code model. To answer this question, I turn now to Dehaene & Cohen's neuronal recycling hypothesis.

3.3 Neuronal Recycling

3.3.1 Introduction to neuronal recycling

The triple code model gives us an idea of how representations of number might be spread out according to the tasks they are required to perform, but it doesn't tell us how the symbolic and linguistic codes manage to link up to the abstract code and allow us to access magnitude information in a way that is precise and allows comparing and understanding numbers of arbitrary size. This job is left for Dehaene & Cohen's (2007) neuronal recycling theory.

For Dehaene, the answer to how we bridge the gap from the content of our evolved cognitive systems to the content of number concepts lies in culture: "Cultural inventions, such as the abacus or Arabic numerals … transformed [the intuition of number] into our fully-fledged capacity for symbolic mathematics" (Dehaene 1997/2011, x). There is a seemingly paradoxical aspect to the importance of culture in the development of numerical cognition. While the fact that innate capabilities like vision and hearing recruit the same cerebral circuitry across cultures and individuals can be explained by the effects of biological selection throughout generations, this does not apply to cultural competences, since these depend on culture-specific learning for their acquisition, and they are too recent to be the result of biological selection pressures,[93] So the question here is, how can we reconcile cultural variability with cortical stability? The answer lies in cultural recycling of evolutionarily-inherited systems:

> To explain this paradoxical cerebral invariance of cultural maps, we propose a neuronal recycling hypothesis, according to which cultural inventions invade evolutionarily older brain circuits and inherit many of their structural constraints. (Dehaene & Cohen 2007, 384)

It could appear reasonable to expect cultural variation to be mirrored at the cortical level. For many, the best way to explain how culture invades parts of our brain is to appeal to a domain-general learning capacity afforded by our brain's unusual plasticity. On this view, the cortex is something like a blank slate on which cultural

---

[93] For example, Dehaene & Cohen point out that writing systems first emerged around 5400 years ago and until very recently only a tiny portion of humanity was able to read. This implies that "it is logically impossible that human brain regions evolved specifically for the purpose of reading" (Dehaene & Cohen 2007, 386).

representations get free reign to reassemble and recruit cortical structures according to their needs.[94] However, such unbridled plasticity fails to explain why the same brain regions would be recruited for specific culturally learned skills like reading and arithmetic: if generalized plasticity were true, then there would seem to be no reason for specific cortical regions to house the variations on cultural invasions, instead of these showing up anywhere within the plastic regions of the brain. Against this general plasticity, Dehaene and Cohen (2007) propose that cultural variability is "tightly constrained by our prior evolution and brain organization" (Dehaene & Cohen 2007, 384). To explain this apparent contradiction in biological stability and cultural variations, Dehaene and Cohen (2007) deploy their *neuronal recycling hypothesis,* according to which culture makes small modifications to pre-existing parts of our brain's architecture whose specific structures makes them ideal candidates for cultural representations like those required to learn reading and arithmetic. In support of this theory, Dehaene and Cohen (2007) present data showing that despite tremendous cultural variation in how members of different cultures read or practice arithmetic, these abilities recruit the same brain regions across cultures – i.e. the hIPS associated with the ANS, for arithmetic, and the left occipito-temporal junction for reading.

The neural recycling hypothesis relies on the following three main postulates:

1) Evolutionarily-driven cortical organization: the main structural divisions and anatomical connections in our brain are drawn by our evolutionary history. These 'neural maps' take shape early on in ontogeny and impose constraints on further learning.

2) Neural niches for cultural features: culturally-transmitted abilities like reading and arithmetic must re-organize parts of our brains. For this to happen, these parts

---

[94] Menary (2015a) often expresses ideas sympathetic to such strong plasticity, as I discuss in chapter 6. Dehaene (2014) disagrees.

of our brain must be sufficiently plastic to allow cultural tweaking, but also already contain relevant structural properties required for specific culturally-inherited abilities. In other words, reading and arithmetic can't recruit any part of our brain, but must recruit on those parts whose architecture can support the development of additional structures required by these tasks.

3) Continuity throughout our lifespan: despite cultural conquering of neural landscapes, the brain's pre-existing structure is not completely erased, and continues to impose constraints on culturally learned skills.

The main predictions made by this model are:

1) Limited cortical variability: cultural abilities should map to the same brain regions across cultures and individuals.

2) Limited cultural variability: there should be invariants across cultures (think Chomsky's minimal project) whose existence can be traced to the neuronal constraints imposed by the cortical areas under cultural invasion.

3) Learning speed should reflect complexity of cultural learning domains: the more distant a cultural representation is from the cortical region it is recruiting, the longer it should take to learn: "Prior cortical constraints should ultimately explain both the ease with which children acquire certain cultural tools and the specific difficulties that they occasionally meet" (2007:385). An example they give is that children have difficulty with letters that are mirror images of each other like 'p' and 'q' because our visual system has a tendency to neglect such differences, since few objects in nature change identity as a result of inverting left-right visual features.

4) Cultural reorganization can incur decrease in certain cognitive abilities: when culture re-wires a brain region, it can make it less efficient at the tasks it had evolved

for. Again, as an example, Dehaene and Cohen mention that reading can decrease symmetry perception.[95]

3.3.2 Support for Neuronal Recycling

Support for the neuronal recycling model of development of numerical abilities comes from brain imaging evidence at different levels of cortical architecture,[96] as well as evidence for evolutionary precursors to culturally-recycled areas, and developmental data concerning the brain areas recruited by cultural practices.

First, imaging studies show variations on cultural practices like reading and arithmetic are associated with cortical invariance at different levels or cortical organisation – or, for those levels where such resolution is not yet possible, the evidence is at least consistent with such invariance.

For example, in the case of reading, the cortical region that Dehaene and Cohen identify as being recycled for reading is a section of the left occipito-temporal sulcus that appears to react preferentially to visually presented letters, earning it the title of Visual Word Form Area (VWFA). This area is recruited across subjects and cultures, despite the incredible variation in how reading symbols present themselves to our visual system (e.g. comparing the alphabet to mandarin or tagalog characters): "Word-

---

[95] For more detail on how this model applies to the case of reading, see Dehaene 2009. Given that our interests lie with arithmetic, we will mostly focus our attention on how neuronal recycling can describe the development of numerical abilities, using the case of reading as an illustration of the way in which pre-existing brain architecture structures cultural invasions.

[96] When they speak of cortical maps, Dehaene and Cohen are describing "a lawful relation between the surface of cortex and a relevant aspect of the representational structure" (Dehaene & Cohen 2007, 385). They describe three possible size scales for these maps: macromaps, which refers to the geometrical relations between brain areas that can span several centimeters; mesomaps, which describes structures within brain areas spanning at most a few centimeters (e.g. mapping from the retina to visual areas); and, more speculatively, micromaps, which are structures typically responsible for individual features (e.g. line orientation in vision) organized at scales of a few hundred microns.

induced activation is found at or around the VWFA site in all good readers, regardless of the writing system they master" (Dehaene and Cohen 2007, 386).[97]

Similarly, despite the ubiquity of Indo-Arabic notation across cultures, there has been tremendous cultural variation in how mathematics and arithmetic have been practiced over the centuries. For example, many south American cultures like the Mayans and others used complex systems of knots to represent and operate over numerical quantities (Ifrah 1998).

As is the case for reading, given that arithmetic is a recent cultural invention, it may be surprising to learn that it reliably recruits the same brain regions across individuals and cultures, despite the variety of cultural practices used to process numerical quantities. It shouldn't come as a surprise that comparing, adding, multiplying, and subtracting with Indo-Arabic numerals reliably recruits the same bilateral parietal areas, given that these tasks all concern the same input format. As mentioned in section 2.3.2, this region appears recruited by any numerical task, irrespective of whether numerical content is presented symbolically or in the form of groups of objects, and irrespective of the modality in which this content is presented (Castelli et al., 2006; Piazza et al., 2004, 2006, 2007; Andres et al. 2012). This same region shows increased activation in mere detection tasks, when compared to the visual and auditory detection of numbers, but not letters and colors (Eger et al., 2003), suggesting that linguistic and non-linguisitc symbols for numbers automatically recruit the magnitude representation of the hIPS,

---

[97]This area does more than merely recognize general visual features like intersecting lines, as evidenced by the fact that its activation is increased for familiar letters when compared with letters from languages unknown to the subject, and from the fact that even within the subject's native tongue, this area responds more to orthographically valid letter strings than random letter strings. Further evidence for this area's specialization can be seen from the fact that it does not respond to upper versus lower case presentation of letter strings. Thus, there is considerable evidence that "the VWFA develops an elaborate functional specialization during reading acquisition – yet this acquisition seems to be remarkably constrained within cortical space" (Dehaene & Cohen 2007:387), despite minor cultural differences in cortical localization or reading areas.

thus confirming the existence of a central, modality-independent semantic representation of number, as predicted by the triple-code model.

As an extreme illustration of the importance of this semantic representation despite cultural variation, there is evidence that even for those cultures where arithmetical practices require motor skills, the hIPS is recruited in calculation. For example, Tang et al. (2006) found bilateral IPS activation in both Chinese and English speakers during calculation and comparison tasks. This study also found cultural variation was reflected in other areas of the brain, with increased activation in the left premotor cortex for Chinese speakers and increased activation of left perisylvian areas for English speakers. These data suggest that mental calculation recruits different networks in subjects from communities whose writing involves more elaborate hand movements or where individuals learn to calculate using an abacus when compared to subjects from communities that learned how to calculate using multiplication tables. The difference in activation patterns likely reflects the difference in external artefacts used when learning how to calculate, showing that our brain has integrated external object manipulation in its circuitry.

These data are consistent with Dehaene's triple-code model, in that the cultural invariance of the bilateral IPS activation reflects the central importance of the abstract representation of numerical quantity recruited in all tasks, while the cultural variability in how this representation is recruited in specific tasks like calculation is reflected in the cerebral variability observed here (see Dehaene and Cohen, 1995). Some cultures rely on language and writing to learn and transmit arithmetical operations, which means that arithmetical facts like addition results and multiplication tables are eventually stored in linguistic format. In contrast, those cultures where writing is more elaborate or where calculation involves manipulating an abacus will recruit motor circuits and arithmetical operations will show increased activation in these areas (see Barner et al. 2016).

A second source of evidence for neuronal recycling is that neuroimaging has identified the structure of evolutionary precursors for culturally novel functions in primate brains, thus showing that the cultural variant recruits a specific, evolutionarily-ancient system due to the function it plays in primate behavior. Cortical invariance in the face of cultural variation poses the question of why a specific brain region is recruited for a specific cultural task. For example, why is the VWFA consistently recruited for reading, rather than other parts of the brain? Neuroimaging of animal brains allows us to answer this question by allowing us to identify what sort of function a brain region area performs in primates. Knowing a brain regions' evolutionary history allows us to identify what role this function plays in cognition in other species, which in turn constrains the possible cultural recycling of its biological function for cultural purposes in humans.

The case of reading illustrates this aspect of neuronal recycling particularly well. There is evidence that an area of macaque brains is hierarchically organized to encode increasingly complex visual feature combinations. The same hierarchical organization appears present in homologous areas of the human brain, in regions that are preferentially activated during reading and letter-recognition tasks (Dehaene 2009). This suggests that the parts of primate brains that have evolved to be good at recognizing and processing specific shapes can be recycled for reading in human brains, since reading requires efficient processing of visual shapes.

While the requirements of reading, including letter recognition, constrain which parts of the brain could be rewired to suit this purpose, the way our brains are organized before culture starts to rewire it constrain possible cultural variations. In other words (pun intended), it looks like written symbols have mimicked shapes characteristically found in nature, thus making them easier to learn and represent, given that parts of our brain have already evolved to specialize in recognizing and processing specific shapes.

Cultural invasions of cortical networks only goes well when these networks were already used to process stimuli similar to those that make up the cultural practice. In the case of reading, the finicky demands of the brain's evolutionarily-shaped architecture – including the anterior-posterior hierarchical organization of the left occipito-temporal sulcus, the left hemisphere's superior ability to process fine-grained stimuli, and the proximity of the VWFA to language centers – constrains the shapes letters can have across languages.[98] Thus, according to Dehaene (2009) letter shapes are not arbitrary, but result from constraints imposed by the brains on what sort of shapes can be efficiently processed. This means that even though we were not born to read, we were born to process stimuli that share sufficient structural similarities with reading-relevant stimuli.

Similarly, while we weren't born to process Indo-Arabic numerals, our brains have evolved to process information about discrete quantities of objects in our environment. Evidence for the evolutionary precursors for arithmetic was presented in chapters one and two, where I summarized findings that strongly suggest that many animal brains are equipped with an ANS, and that there is evidence that it is located in homologous areas in primate brains (section 2.4). The fact that macaque brains show similar structural organization of activation for attention, eye saccades, and hand-related actions like grabbing to the layout found in humans by Simon et al. (2002) suggests that there would be precursors to numerical abilities in macaque IPS. This finding has been confirmed by Nieder and colleagues (e.g. Nieder & Miller 2004), who found individual neurons that displayed preferential activation for specific numerosities going beyond 30 (Nieder & Merten 2007), as mentioned in chapter 2. This offers evidence

---

[98] Dehaene and Cohen offer three main lines of evidence to explain that the VWFA is constrained to the same cortical location across cultures. First, the posterior-anterior organization of increasingly complex feature combinations has precursors in primate brains (Rolls 2000) and has been found in human image recognition using fMRI (Lerner et al., 2001). Second, its lateralized location can be due to the fact that this location typically represents foveal stimuli, whose informationally loaded details requires fine-grained high-fidelity processing to capture. Third, its location in the left hemisphere could be explained by the fact that this offers shorter connections to language areas than a right lateralization, or by the fact that the left hemisphere is typically superior for fine-grained processing required for fast letter recognition (Kitterle and Selig, 1991).

that the functional localization of human numerical abilities observed in cortical maps was constrained by pre-existing structures containing neurons that code for numerosity.

The data showing the presence of an ANS in areas of primate brains homologous with the hIPS in human brains explains why the hIPS is recruited for arithmetical tasks, instead of other parts of the brain: given that the role of the ANS is to track quantities of discrete objects, it is particularly well suited for numerical cognition, which involves a more precise treatment of precise quantities. As I discuss in more detail in chapter 6, those cultural practices that mark a better fit with the ANS will have greater chance of surviving, given that they require less effort to master: "the human brain is endowed with an innate mechanism for apprehending numerical quantities, one that is inherited from our evolutionary past and that guides the acquisition of mathematics" (Dehaene 2011, 30).

We saw in chapter one that infants and animals are sensitive to changes in numerical information both inside and outside the subitizing range, as demonstrated in a variety of violation-of-expectancy paradigms (section 1.3.3). We also saw that the same brain regions in the bilateral intraparietal sulci are active in numerical tasks in infants and adults, and that homologous brain regions in animals are recruited in numerical tasks. The data summarized in the first chapters showed that "results suggest that availability of a functional parietal quantity system is an essential pre-requisite for arithmetic development" (Dehaene & Cohen 2007, 392).

Strong validation of the neuronal recycling hypothesis would come from evidence that culturally-driven learning changes the structure of the hIPS in order to make representations of natural numbers possible for humans. Perhaps due to the resolution required to detect such culturally-driven changes, and the ethical and practical constraints on methods we can use to probe human brains, so far, no direct evidence has been discovered for or against the claim that cultural practices change the ANS. However, some findings nevertheless do provide indirect evidence. For example, the intraparietal sulcus appears to be increasingly active in numerical tasks as people age

(Ansari and Dhital 2006; Rivera et al. 2005), which hints at increased development of this area in relation to advancing arithmetical training.

Another source of evidence for recycling comes from the previously mentioned correlation between arithmetical skill and ANS acuity: the sharper a child's ability to distinguish between stimuli based on their numerosity, the better the odds are that the child will perform well in arithmetic (Gilmore et al. 2007). Given that arithmetic requires the ability to manipulate symbols for quantities, if an individual's ANS somehow performs better than average, then that person may find it easier to recruit the ANS for symbolic use. An extreme illustration of this can be found in children with dyscalculia, who often show decreased activity in parietal areas, thus making the relation between symbols and the ANS less efficient.

3.3.3 Neuronal Recycling and mapping to the ANS

Such evidence suggests that advanced arithmetic and its many cultural variations is based on a mapping of cultural practices for quantification on the quantity representation provided by the ANS. This explains why different cultural practices recruit the same brain region: this is where our representation of numerical magnitude lies, and this representation is needed in order for the symbols and artefacts used in arithmetical practices to be meaningful:

> 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. We would then be reduced to purely formal manipulations of digital symbols, in exactly the same way that a computer follows an algorithm without ever understanding its meaning. (Dehaene 2011, 75)

Thus, the hypothesis of neuronal recycling explains how the semantic representation of numerical magnitude housed in the hIPS gets exploited by cultural practices by a form of mapping from symbols and practices to magnitude representations in the hIPS:

"our understanding of the cultural symbols of numbers is grounded in links with neurons coding for specific nonsymbolic numerosities in intraparietal cortex" (Dehaene & Cohen 2007, 391).

I have already mentioned the evidence for the fact that there is a mapping between the ANS and numerals in chapter 2, where I showed that its signature distance and size effects are present in numerical tasks involving numerals and other symbols for numbers. The fact that behavior in symbolic tasks displays scalar variability that follows from Weber's law is a clear indicator that numerals are eventually mapped to the ANS, which grounds these symbols and gives them numerical content.

Whether a culture uses words, non-linguistic symbols, abacuses, or tallies made from notches on bones to count things, such numeration practices recruit the same semantic representation of numerical magnitude to orient behavior. The reason why there is cortical stability despite cultural variation in the practice of arithmetic is that the ANS has evolved to track quantities, which makes it the perfect – and only – candidate to give meaning to the more precise quantification tasks that make up the culturally variable practice of arithmetic. On this model, the ANS acts as a general representation of what quantities of discrete objects are, and cultural quantificational practices recruit this representation and sharpen its approximate output by providing labels for specific representations of quantities.

This means that, according to Dehaene, our more advanced numerical skills require the use of external symbols for numbers:

> How did *Homo sapiens* alone ever move beyond approximation? The uniquely human ability to devise symbolic numeration systems was probably the most crucial factor. Certain structures of the human brain that are still far from understood enable us to use any arbitrary symbol, be it a spoken word, a gesture, or a shape on paper, as a vehicle for a mental representation. Linguistic symbols parse the world into discrete categories. Hence, they allow us to refer to precise numbers and to separate them categorically from their closest neighbors. Without symbols, we might not discriminate 8 from 9. (Dehaene 2011,79)

The basic idea here is that the fact that humans use symbols allows us to have precise representations of otherwise entangled stimuli. This idea, according to which moving beyond the limitations of our evolutionarily-inherited numerical abilities necessitates external symbols, is an almost universal posit in the numerical cognition literature. We will see another version of it in the next chapter, which summarizes Susan Carey's externalism, as well as in other externalist accounts presented in later chapters. Before getting to these other externalist accounts, however, it is important to highlight some of the shortcomings of Dehaene's neuronal recycling hypothesis and the associated claim that culturally-induced mapping recycles the circuitry of the ANS to allow advanced numerical cognition to happen.

3.4 Problems with Dehaene's mapping

I have just summarized Dehaene's triple-code model of the cognitive architecture underlying advanced numerical cognition and some of the reasons motivating its main properties. I also summarized the main lines of Dehaene and Cohen's neuronal recycling hypothesis and how it describes the development of an understanding of what natural numbers are by a mapping from culturally-acquired symbols for numbers to representations produced by the ANS, and that this mapping re-wires the ANS and sharpens its representations.

One problem with his account is that it fails to provide an explanation of how the mapping occurs: while we can agree that the triple-code model describes the neuronal architecture underlying our numerical abilities, this is essentially a description of the end result of a learning process that we need to account for if we want to have an understanding of the acquisition of number concepts. To an extent, the neuronal recycling hypothesis also seems to take the existence of developed numerical content as a *fait accompli*, in that it describes a potential change in neurological structure of the brain that could allow the development of numerical content, but fails to provide a learning mechanism that initiates or allows this learning to take place. In other words, while we can agree with Dehaene that cultural practices like reading and advanced arithmetic rewire parts of our brain that are especially suited for these practices, this does not help explain how the mapping between culturally-learned symbols and the ANS can help learn what natural numbers are, given that neither the symbols nor the ANS embody some of the main characteristics of natural numbers, such as their precision and their infinitely extendable domain.

As an illustration of the explanatory limits of Dehaene's mapping, consider Susan Carey's criticism of this mapping as a potential learning device capable of explaining the development of natural number concepts. Carey's criticism of the mapping-as-learning mechanism proposed by Dehaene is mainly based on two empirical challenges for any theory of the development of natural number concepts. The first challenge concerns "how the child creates the initial mappings that get the process started" (Carey 2009, 313), given data showing that this initial mapping requires a stepwise and lengthy learning period. The second challenge is that there is evidence that children learn the meaning of many number words before this mapping is complete.

3.4.1 Mapping in stages

To see what the first problem is, consider Wynn's (1990, 1992b) experiments using the Give-A-Number tasks, where children are asked to give the experimenter a precise number of objects. These studies show that children reliably go through the same stages in learning the meaning of number words. For English, these stages go as follows: The No-Numeral-Knower stage describes when children are unable to give even one object when asked to. Then, between 24 and 30 months of age, they correctly give one object when asked to, but fail at any other number, at which point they are One-knowers. 6 to 9 months later, they become two-knowers, where they can correctly give one object when asked to, or two objects when asked for two, but give random numbers of objects for any number word larger than two. Another, distinct three-knower stage follows. Carey calls children at these stages 'subset-knowers' to highlight the fact that, for 12 to 18 months, they only know how to correctly apply an initial segment of the count list.

The crucial knowledge stage almost invariably occurs once children have become three-knowers (in some rare cases, after being four-knowers):

> around age 3 1/2 on average, English middle-class children become cardinal principle knowers—they work out the numerical meaning of the activity of counting and can now reliably produce sets with the cardinal value of any numeral in their count list. (Carey 2009, 298)

Once children become Cardinal-Principle-knowers (CP-Knowers hereafter), "they have created a representation of some positive integers, a numerical representations that transcends core number representations" (Carey 2009, 302). According to Carey, considerable developmental evidence shows that these stages occur at roughly the same ages and that the qualitative shift between being a subset-knower and being a CP-knower is found in a number of different tasks that measure numerical abilities, suggesting within-child consistency and regular learning milestones.

As an example of the qualitative behavioral shift between CP-knowers and subset-knowers, consider that in Give-a-Number studies (Wynn 1990, 1992a), subset-knowers do not count items when asked to produce a specific number of these, whereas CP-knowers do. When asked to count items and then identify their number, subset-knowers will correctly apply the counting routine to items, but fail to produce the right number word, while CP-knowers can both apply the labels correctly and invariably use the last number word named as the number of items to be counted.[99] This suggests that subset-

[99] The same distinction was observed in Whats-on-this-Card studies, where children are asked to describe the objects shown on cards, including their number (LeCorre et al. 2006). In other words, subset-knowers on Give-a-number tasks are also subset-knowers on the What-On-This-Card task, despite the considerable differences in processing requirements between these.

knowers fail to understand the cardinality principle, according to which the number of items in a collection is labelled by the last number word used to count these. The basic difference is that subset-knowers' answer to 'how many?' questions do not match the last number in their count list. According to Carey, this shows that "many different tasks provide evidence for a qualitative change in understanding counting upon becoming a cardinal principle knower" (Carey 2009, 300), since children are subset-knowers or CP-knowers across many different numerical tasks, and even within the subset-knower stages, the limitations associated with each stage are observable in both tasks mentioned above.

According to Carey, there is no reason why a mapping from lists of number words to ANS representations like that proposed by Dehaene would go through such a protracted, piecemeal learning curve, since the ANS is not structured this way. For example, there is no singular-plural distinction in the ANS, and there appears to be no important distinction between the ANS representation for '1' and those for the next number words in a count list. This then fails to explain why children first learn the meaning of 'one' and only much later learn the meaning of other number words. Unless there is an explanation for these stages of learning, the mapping-based account does look less appealing.

3.4.2 Mapping *after* understanding

The second empirical challenge for a mapping-based learning mechanism is that subset-knowers must have a partial mapping of number words to the ANS before they become CP-knowers, since the transition involves an induction on how the first number words are associated with content from the ANS, so that NEXT NUMBER WORD comes to be associated with LARGER ANALOG MAGNITUDE. And yet, according to Carey, "Children apparently integrate numerals with analog magnitude representations some six months *after* they have learned how counting represents number" (Carey 2009, 314). Thus, more than being a mere logical possibility, it looks like mapping number words to representations of the ANS takes place after children have figured out the meaning of the words in their lists of number words.

As evidence for this claim, consider the fact that no evidence has been found that subset-knowers have partial mappings. Researchers have sought to find evidence that subset-knowers have partial mappings between counting lists and ANS representations, since this would show that such a mapping is involved in the gradual learning of the meaning of number words. To do this, they have tried to show that subset-knowers can estimate the number of objects in collections with numerosities outside the subitizing range, since this would preclude possible influence of natural language quantifier representations or parallel individuation. For example, Condry & Spelke (2008) showed children at different stages of subset-knowing cards displaying either four or eight objects, and found their answers to be at chance levels. According to Carey, this means that the number words had not yet been mapped to the ANS, since the ANS can distinguish between these quantities. Similarly, LeCorre and Carey (2007) quickly flashed cards displaying 1-10 images and asked children of varying subset-knowing levels to estimate how many were on the card and found their answers to be at chance.[100]

According to Carey this shows there is no evidence that subset-knowers have mapped any numerals onto representations of the ANS, nor that they tend to use larger number words for larger quantities, which suggests that the understanding they have of number words isn't based on a mapping, given that their use of number words does not mirror the structure of the ANS. Even recently-christened CP-knowers appear to lack this mapping, as illustrated by the fact that children can use counting routines to determine the number of objects in a collection but still fail to use larger number words for larger

---

[100] See also Lipton & Spelke 2005.

quantities when asked to identify the number of objects on briefly flashed cards (LeCorre & Carey 2007). When LeCorre tested 3- to 5- year old children that were CP-Knowers, he found that younger children (average age 4:1) did not use larger number words for larger numerals outside the subitizing range, but that older children (average age 4:6) did.

Carey points out that the fact that we go through levels of subset-knowing one after the other in distinct stages before becoming CP-knowers shows that the ANS is not involved in the early stages of the learning process, given that the first few numbers whose meaning we learn are in the subitizing range, not the ANS's range. According to Carey, "These data absolutely rule out the possibility that mapping numerals in the range of 5 to 10 to analog magnitudes plays any role in the construction of the numeral list representation of natural number." (Carey 2009,316). Rather, "Numerical meanings of "one," "two", "three" and "four" and nothing more underlie the construction of the counting principles" (Carey 2009, 317).

Two final empirical problems that are more specific to the ANS mapping theory are worth mentioning: one, on this model, "the very striking discontinuity at "four" has no ready explanation" (Carey 2009, 318), since there is no marked division at four in the ANS, so that there is no reason why partial mappings would go through subset-knower stages rather than having mapped up to 'six' or 'ten', for example.

Last, and perhaps more importantly. Carey points out that children's estimates of quantities in the subitizing range do not display scalar variability, while it does outside this range. This means that the ANS cannot underlie behavior here since its main signature does not describe behavioral regularities in the range in which the partial mapping would take place. Rather, according to Carey, children's estimation abilities here suggest that parallel individuation explains behavior in this range.

As Carey points out, however, even when children become CP-knowers, their representations of numbers usually do not go beyond 20 (Carey 2009, 335), which means that further associations from larger numerals to numerical content need to be built. According to Carey, this is where a mapping to the ANS becomes useful – and possible. This means that, regardless of whether or not we agree with Carey's arguments against an ANS-based mapping for learning the meaning of the first few numerals, such a mapping can play a part later in the development of number concepts.

In this chapter, we saw that while empirical support for Dehaene's triple-code model of numerical cognition is strong, his proposed solution to the gap problem, which is based on the neuronal recycling of our brain by a mapping from number words and symbols to the representations of the ANS, does not seem to fit well with the available developmental data concerning how children learn the meaning of the first few number words.

This discussion of Carey's arguments against Dehaene's proposal that a mapping of culturally-acquired symbols onto the representations of the ANS can explain the origins of our understanding of natural numbers highlights an important lesson for any potential explanation of the origins of numerical cognition: not only should any such account pass the empirical tests imposed by data on subset-knowers, but it is important not to confuse an description of the cognitive architecture behind our numerical abilities for an explanation of how we manage to learn what natural numbers are. Given that Dehaene's mapping seems well supported empirically, that the existence of such a mapping seems established beyond doubt, and that there is evidence that there is an understanding of what natural numbers are before this mapping – or at least, there is an understanding of what an important initial segment of these is – any account of the development of advanced numerical cognition owes us an explanation of how we come to realize that the last numeral in a count list labels a precise quantity of discrete objects. Dehaene does not appear to have such an explanation. Carey, however, does. In the next chapter, I summarize the main lines of Carey's explanation of the learning device she claims is responsible for our understanding of what natural numbers are.

CHAPTER IV

CAREY'S BOOTSTRAP

4.1 Introduction: the problem of learning

The question of how knowledge is acquired has puzzled philosophers since the pre-Socratics. The problem can be framed along Fodorian lines as follows: if learning involves confirming hypotheses, how can we learn anything, given that in order to confirm a hypothesis, one must already know what the hypothesis is about? In other words, we can't learn what we can't represent, and we already represent what we know. While Socrates' reliance on reincarnation to explain to Meno how the slave managed to recognize Pythagoras' theorem finds few adherents these days, the nativist explanation of the origin of our knowledge of the world is alive and well, despite the association of science with empiricism.

In her 2009 magnum opus, *The Origin of Concepts* (TOOC hereafter), Susan Carey carves a nativist narrative out of developmental and animal studies to tackle the age old question of how learning is possible. To show that the tools and methods of modern science have allowed us to make progress in answering questions concerning innate knowledge, Carey has at her disposal the rewards of a lifetime of work dedicated to studying the human conceptual apparatus, which she uses to piece together a theory of learning that explains how it is possible for people to acquire representations whose content is discontinuous with the representations they are built from, as explained below. To make her case, Carey weaves together scores of empirical studies into a coherent, philosophically compelling account of what concepts are and how we learn them. One of the concepts she discusses is detail is the development of the concept NATURAL NUMBER.

To understand how Carey explains the development of NATURAL NUMBER, it is essential to understand her account of the development of the human conceptual apparatus. This is because Carey's account, like Dehaene's and many others, relies on innate representations with numerical content to explain where natural number representations get their numerical content from. If we want to understand how innate representations can have conceptual content of the sort that could serve as the basis for advanced numerical cognition, Carey's core cognition framework arguably offers the best explanation, since it attributes a special status to representational systems like the ANS and the OFS in its taxonomy of cognitive systems. In this chapter, I walk us through Carey's account of the ontogenetic development of numerical content. Section 4.2 sketches Carey's distinction between perception and conception, in order to set up her account of core cognition in section 4.3. I then discuss evidence for core cognition in humans in section 4.4 before summarizing Carey's evidence for a third representational system with numerical content, which Carey calls *set-based quantification.* This notion will then be used to explain how we learn through conceptual discontinuities (section 4.6) via Carey's account of learning in terms of what she calls Quinian Bootstrapping (section 4.7). I close the chapter by setting up my origins problem and applying it to Carey's account in section 4.8.

4.2 Perception vs conception

Carey endorses a representational theory of mind in which "representations are states of the nervous system that have content, that refer to concrete or abstract entities (or even fictional entities), properties, and events" (Carey 2009, 5).[101] Carey's inventory of the contents of our minds includes perceptual/sensory[102] representations and conceptual representations, as usual. Less common is her inclusion of representations that share properties of both conceptual and perceptual representations, which are the representations of core cognition. To understand how core concepts differ from other representations, we must first understand how Carey distinguishes between concepts and percepts.

On Carey's view, concepts differ from perceptual representations in two important ways: first, concepts figure into different computations and inferences than perceptual representations; second, the mechanisms that fix conceptual content differ from those of other types of representations. Consider perceptual representations first: these connect us to the outside world, our senses translating and transducing the way our environment interacts with our body into nervous signals that find their way to the brain, where they act as sensory and perceptual representations of the parts of the world evolution has equipped us to react to. Perceptual representations get their content from things in the world, while the mechanism responsible for this causal content-fixing can be explained by biological evolution. Representations of round things are about round things, because we have evolved to recognize shapes of things in the world. Importantly, this means that perceptual representations are largely innate, in that they do not require any learning to be acquired.[103] At most, they require a form of triggering

---

[101] While Carey does not offer any explicit definition of what mental content is, she treats this notion as being equivalent to the domain of a representation, or the set of all things to which it applies. Contrast this with Cohen Kadosh & Walsh's definition of representation, according to which representations are "patterns of activation within the brain that correspond to aspects of the external environment" (Cohen Kadosh & Walsh 2009, 314).

[102] For the most part, Carey does not bother distinguishing between perceptual and sensory representations, since this distinction does not have much impact on her theory of conceptual development. I also adopt this practice throughout the text.

[103] It is important to realize here that for Carey, 'innate" simply means 'not learned'. What is innate need not be present at birth, nor does it mean that we have innate representations without our environment

from the environment.

Such perceptual representations are limited in their inferential role. For example, on its own, the representation of the color, temperature, or shape of an object does not tell us much, since the mere presence of a percept does not give us information about the identity of the object that bears that percept. In this sense, perceptual representations are limited to what is present to our mind at a certain moment. Further limiting their informational role is the fact that perceptual representations are often thought to be informationally encapsulated, since they are insensitive to information coming from other representational systems.[104]

In contrast with these perceptual representations, concepts can take part in much more complex and elaborate inferences. Concepts are described as being informationally promiscuous, in that they will take input from many different sources and participate in domain-general computations involving many modalities. A representation with the content RED can only figure into simple modality specific inferences such as NOT YELLOW. In contrast, upon hearing a certain sound, my concept WOLF allows me to infer that an animal with four legs and fur and teeth is nearby, which means that I might want to escape. Thus, perceptual representations have limited and impoverished

---

first tokening these via the senses: "claiming that representations of red or round are innate does not require that the child have some mental representation of red or round in the absence of experience with red or round things" (Carey 2009, 12). There are many reasons to think this is true. For example, as Carey points out, while stereoscopic depth perception is not present at birth, it invariably develops at around six months of age.

[104] Jerry Fodor (1983) famously came up with a two-part distinction about mental representations based on their inferential role. The notion of informational encapsulation is meant to describe the fact that some cognitive processes are immune to input from other systems. A classic example used by Fodor is that of the Müller-Lyer illusion: despite knowing that the lines are of the same length, our visual system will still process the information the same way and thus they will continue to look like one is longer than the other. More central representations figure in more inferences and can thus figure into theories made up of complex interrelations between conceptual representations.

inferential roles when compared to those of concepts.

Carey also distinguishes concepts from percepts in terms of how each type of representation gets their content. Unlike perceptual representations, conceptual content is often fixed via social learning and by involvement in inferences that often involve other conceptual representations with no direct link with the world outside our head. Thus, to understand how concepts get their content, it is not enough to look at the things that caused their tokenings, since the mechanisms that fix the content of concepts are not limited to causal determination by external objects. Rather, Carey adopts Block's (1986, 1987) two-factor theory of conceptual content, where a concept's content is determined by its inferential role and by the objects that caused its tokening (Carey 2009, 5). Lexical concepts like word meanings are typical exemplars of concepts in Carey's way of slicing up our mind. The content associated with a word is determined by the objects to which the word applies as well as the inferences warranted by this content. Thus, the word 'wolf' is associated with wolves, but also with socially and culturally acquired content that can vary from person to person, such as FEARED BY GRANNY.

4.3 Core cognition vs other representational systems

While there is nothing particularly controversial about these distinctions between perceptual and conceptual representations, Carey's representational repertoire also includes less common conceptual representations that are the output of innate, modular perceptual analyzers: these are the representations of core cognition. Adapting elements of Elizabeth Spelke's work on core knowledge (e.g. Spelke 2000),[105] Carey

[105] Carey prefers the label 'core cognition' to Spelke's 'core knowledge' because "the representations in core cognition need not be (and often are not) veridical and therefore need not be knowledge." (Carey

argues that representations like AGENT, OBJECT, and NUMBER are the output of distinct cognitive systems that share features of both perception and conception. Providing proof of the existence of core cognition is one TOOC's main theses.

To demonstrate the existence of such innate, domain-specific conceptual representations, Carey discusses two beautiful examples taken from the animal kingdom: the representation of the night sky in indigo buntings, and the imprinting behavior of newborn chicks. Consider first the buntings. These migratory birds are born with a domain-specific learning device that allows them to compute where north is, thereby allowing them to fly in the right direction when seasonally-induced hormonal changes signal them that it is time to migrate. By creating artificial night skies in a planetarium, researchers showed that these birds would orient their flight patterns towards the center of rotation of the sky, irrespective of configuration of the stars or any other feature of the sky (Emlen 1975).

It thus appears that these birds are equipped with an innate perceptual input analyzer that takes as input a visual representation of the night sky and from this input computes the location of its center of rotation - which is where the north star lies. The reason why this illustrates what core cognition is all about is that the representational system in charge of producing a representation of the center of rotation of the night sky is neither purely perceptual nor completely conceptual. It is perceptually encapsulated and domain specific, in that it takes only input from the visual domain. And yet, the content produced by this system is not merely perceptual, since the information about the azimuth of the night sky does not only present aspects of the world as they are in the here and now: no description of the night sky's perceptual attributes, such as the color of the sky, of the stars, their location and configuration, would include the content NORTH, or IMPORTANT LOCATION. Rather, a representation of the center of

---

2009,10)

rotation of the night sky takes many perceptual representations and builds a more complex one that is not immediately perceivable, and that is available to other systems and modalities (e.g. those in charge of navigation). However, given that this representation of north only plays a role in a limited set of further computations, it is not fully conceptual.

Another good illustration of a core conceptual module is observable in the behavior of newborn chicks. In famous experiments, Lorenz (1937) showed that these animals are equipped with two innate learning devices that evolved to maximize the chick's chances of closely following their mother soon after birth. One such 'imprinting' module is tasked with detecting certain types of motion and getting close to the thing that moves in hen-like fashion. In case this fails, another system of perceptual input analyzers is charged with detecting certain shapes in the environment - those that somewhat resembles an adult chicken - and again tells the chick to follow the object that looks this way (Johnson et al. 1985).

Both of these cases shows that there appears to be something more than mere sensation in the innate representational systems, and yet also something not-quite-conceptual, thus strongly supporting Carey's contention that core cognition is a representational category of its own, sharing characteristics of both perception and cognition. More specifically, Carey lists six features of core cognition, one of which is shared with conceptual representations, and five that are of a more perceptual nature.

Much like conceptual representations, representations in core cognition have rich inferential roles: the output of core cognition modules are inferentially promiscuous and available for general processing. Also, as illustrated in the cases of the indigo buntings and chicks, their content cannot be expressed in purely sensorimotor or perceptual terms, which is limited to presenting information about the perceptual here and now. While chicks can react to specific shapes and behavior, the mere presence of a shape or specific type of motion does not provide any further information, on its own, nor does the presence of stars in certain configurations in the sky. Rather, for these arrangements of matter to mean something important for an animal, a specific representational system must analyze the perceptual input and produce a representation with the content NORTH or HEN. This being said, core concepts do not participate in as rich an inferential life as explicit conceptual representations. Also, while the content of core cognition is mainly determined by causal interaction between perceptual input analyzers and the environment, the content/extension of fully-fledged concepts can also be determined by socio-cultural processes. While core cognition is the result of innate domain-specific modules, regular concepts need not be. Thus, "core cognition is conceptual, but not fully so" (Carey 2009,15).

Core cognition shares five main characteristics with perception. First, like perceptual representations, the representations of core cognition are the output of innate perceptual input analyzers. This means that we can explain the origin of the content of core cognition modules by appealing to their evolutionary origins:

> Natural selection has constructed these analyzers specifically for the purpose of representing certain classes of entities in the world, and this ensures that that there are causal connections between these real-world entities and the representations of core cognition. (Carey 2009, 67)

Thus, unlike most concepts, whose origins can be traced to social learning and complex inferences, the content of core cognition modules is determined by the things in the world we have evolved to track:

> The requisite causal connection between entities in the world and symbols in the head is guaranteed by natural selection. Thus, we have the beginnings of a theory of how the symbols in core cognition come to have the content they do. (Carey 2009, 116)

The evolutionary origins of core cognition explain another characteristic of these, that some of our core concepts are shared with other animals. This again does not seem to apply to most regular concepts: presumably, no animal has the concept CAR, let alone

ELECTRON.

Third, core cognition modules are in operation throughout our lives. This means that even though core cognition representations are learning devices that serve as the foundation for more advanced concepts and theories, the content they provide is not replaced by more advanced concepts. This is not true for run-of-the-mill concepts, which can be entirely replaced by competing representations. For example, the content of theories about what matter is or reality are can be overturned by scientific learning or religious conversion. This is not true of core cognition.

Fourth, core cognition modules are domain specific, in that they do not take input from any cognitive system or domain. For example, indigo bunting's representation of NORTH only takes as input visual information. This is not true for fully-fledged conception, since concepts can be the output and input of many cognitive domains. For example, words can be presented in visual, auditory, and tactile domains.

Last, and more controversially, Carey claims that the core cognition representations are iconic, or analog, which means that their structures somehow mirror those of the things they represent. Carey gives the example of a picture of a dog: parts of that picture represent parts of the dog. This is not the case for the word 'dog', since a part of that word, say 'og', or 'g', does not represent a part of a dog.[106]

In sum, core cognition is somewhere between perception and full-blown conception. Given the enormity of the body of ethological literature, there is good reason to believe that innate representations with conceptual content in animals are more the norm than the exception. This means we can expect to find cases of innate representations with

---

[106] Whether this distinction between iconic and symbolic format applies to core cognition is controversial (see Ball 2016). Thankfully, it is not one of the most distinctive aspects of core cognition, and even if it were false that core cognition was iconic in format, this would not mean that there is no core cognition.

conceptual content in humans. This is indeed the case. In the next sub-section I show that the ANS and OFS meet Carey's criteria for core cognition, and thus that they can be considered as sharing properties of both perception and cognition. Framing these systems as examples of core cognition is important, since it gives an evolutionary origin to our numerical abilities, which in turn allows us to sidestep possible complications associated with explaining how children come to construct representations with numerical content from purely perceptual building blocks.

4.4 Core Cognition in humans

Like the indigo buntings' innate ability to compute the center of rotation of the night sky, the newborn chick's ability to detect bird-shaped objects or bird-like motion, humans are also born with many innate perceptual input analyzers. Perhaps one of the best known examples of these are human face-detectors, which allow newborns to react to visual stimuli with facial characteristics.

We can easily see why Carey sees the ANS and the OFS as examples of human core cognition modules. First, neither of these produces content that can be expressed in purely spatiotemporal vocabulary. For example, as I showed in section 1.3.4, object files represent objects as being spatiotemporally continuous, despite occlusion. While this may seem like a minor point, it does illustrate the limitations of perceptual systems: the fact that redness and roundness are represented at one instant and are no longer represented at another instant does not, on its own, say anything about the object that was red and round. The fact that object files individuate objects as persisting despite occlusion shows that their content is above and beyond purely perceptual data. As for the ANS, the fact that there are objects in various places in our visual field does not, on its own, give us any quantitative information on these objects, which means that the representations of the ANS cannot be expressed in purely spatiotemporal terms either.

As for how these systems embody the other characteristics of core cognition, both can produce content that can figure into more general computations, as seen by the fact that they can guide action (e.g. reaching for food) and that both can figure into multimodal representations, as illustrated by the work of Meck & Church (1983, section 2.2.5) for the ANS and Starkey & Cooper (1980, section 1.3.5) for the OFS. As seen earlier, both of these cognitive systems are present in other animal species, which suggests that they are evolutionarily ancient. Both also co-vary with the outside world, tracking objects, or quantities of discrete objects. Both also have specific domains of objects to which they respond: object indexing in MOT privileges spatio-temporal information, which means that object files themselves are informationally encapsulated, to an extent, from property and kind information. The OFS also fails to individuate continuous substances like sand and water, while the ANS fails to respond to non-numerical information. Both also stay active throughout our life, as seen by the fact that the subitizing range is the same for adults and children, and by the fact that distance and size effects are present even when numerical magnitude is presented symbolically.

As for format, we saw in chapter one that the format of the ANS appears to be akin to a mental number line, where quantities of objects are represented as positions on this mental number line. If this is true, then the format of the ANS is indeed iconic, since the number 8, which could be represented as ________, does contain the number three, which could be represented as ___. The same holds regardless of whether the actual implementation of this number line is in terms of increased activation by summation units, or increase in the number of neurons activated (see Nieder & Merten 2007; Nieder & Dehaene 2009). As for the OFS, each object individuated opens a corresponding object file, so the format of its representations is iconic, in the sense that each part of the output of the OFS represents a part of the collection of objects it is representing.

In sum, Carey offers compelling evidence that the OFS and ANS display the main characteristics of core cognition. This is important for our purposes because the fact that these systems share properties of perception, on the one hand, explains the origins of the content they produce. We can therefore explain that we have representations of objects and of quantities by appealing to the selective pressures that could have forged representational systems with such content in animals by the increased fitness it bestows on organisms that have such systems. The fact that core cognition modules are evolutionarily ancient means that the origins of some concepts can be explained by the selective advantage our ancestors had of tracking certain regularities in the world, as is the case for perceptual representations: "the extension of the symbols that articulate core cognition is fixed, in part, by evolutionarily underwritten causal relations between entities in the world and representations in the mind" (Carey 2009, 11). This means we do not need to appeal to a learning mechanisms that would operate over our experience of the world, as Piaget and other empiricists would have it, to explain the origins of some concepts.

The fact that these systems share properties of cognition, however, also allows us to explain their general applicability and their ability to figure as inputs to cross-modal and general cognitive processes. This is a welcome result since it can help explain the generality and abstractness of numbers: if it turns out that we can explain the origins of natural numbers by the operation of a kind of learning mechanism over the content of core cognition modules, we can appeal to the generality of these modules to explain why it is that we can apply numbers to any discrete objects of thought, since it is one of the main properties of core cognition modules that they can figure in such domain general inferences. Since the output of core cognition modules is available for more general processing, they allow us to build up more complex representational content using core cognition as its basis. Thus, " The core cognition hypothesis provides part of the solution to our quest for the origin of human concepts, for it consists of systems of innate conceptual primitives" (Carey 2009, 69). The unique characteristics of the representations of core cognition - the fact that these are both innate and yet have conceptual content - means that they can figure into explanations of how it is possible to develop conceptual content from innate perceptual analyzers. According to Carey, "core cognition is the developmental foundation of human conceptual understanding" (Carey 2009, 11).

Before taking a look at how Carey wishes to exploit the conceptual content of core cognition modules to explain how we bridge the gap, however, there is another core cognition module that needs to be examined. This is because Carey, unlike Dehaene (2011), Wynn (1998) and Gallistel & Gelman (1992), does not think that the ANS plays a major role in bridging the gap. Rather, Carey's account of how we bridge the gap relies on the OFS, as well as our ability to put objects into one-to-one correspondence, but also another core cognition module with quantity-related content. In the following section, I outline the reasons motivating this choice of building blocks for Carey's gap-bridging story and introduce the properties of the extra core cognition modules that are present in Carey's account. This involves taking a closer look at what sort of role language plays in this process, if any.

4.5 Set-Based Quantification

It would appear natural and intuitive to presume, as many have in the past, that our numerical abilities derive from our language faculty. Language seems like the perfect candidate to explain the origins of natural numbers, given that it allows endless recursive production and recombination of words, and that such recursive construction out of previously available representations seems like it could describe how we produce new representations of numbers via recursive recombination of previously constructed numbers. And yet, given the evidence presented in chapters 1 and 2, we can easily rule out any account of the development of numerical cognition that would base itself solely on the resources of natural language, given the central role played by language-independent systems like the ANS and the OFS in our numerical abilities.

This being said, while the evidence of their presence in animals and preverbal infants rules out that the ANS and the OFS are based in our language faculty, Carey presents compelling evidence suggesting that the quantificational resources used by natural language quantifiers are not the product of these core cognition modules. If this is the case, then there could perhaps be a role left for language after all, since our ability to think about natural numbers could be constructed from the quantificational resources of natural language, in conjunction with the ANS and OFS. However, as I describe in this section, Carey offers compelling evidence that natural language quantification itself is based on a language-independent core cognition module that processes separate quantity-related information from that produced by the ANS and the OFS. To anticipate, the reason why this matters for our gap problem is that this third system plays a crucial role in Carey's explanation of how we bridge the gap, as I detail in the next section.

When speaking of natural language qualifiers, the first distinction that comes to mind is the singular-plural distinction, which highlights the difference between collections of objects containing exactly one member and collections that contain any other non-zero number of objects. Other natural language qualifiers like 'some', 'many' and 'both', for example, also allow us to track quantities of objects in our environment by giving us information about how many objects fall under a predicate. Given the important role of the ANS and the OFS in our numerical abilities, a natural question here is whether we can explain the quantitative content of natural language quantifiers in terms of the output of the ANS and the OFS. Carey provides two sources of evidence that this is not the case:

First, natural languages include explicit symbols for quantifiers that are not

> represented in the parallel individuation [i.e. OFS][107] and analog magnitude [i.e. ANS] quantificational systems. Second, natural languages deploy concepts for individuals other than spatio-temporally defined objects. (Carey 2009, 254)

The first claim here is that neither the OFS nor the ANS can produce *explicit*[108] representations whose content could be synonymous with that of natural language quantifiers like 'some', 'all', 'none', 'many', or 'a'. This is because the OFS only produces explicit representations of objects, while the ANS only produces explicit symbols for approximate quantities, whereas natural language quantifiers produce explicit symbols for quantity-based selections of members of collections. For example,

---

[107] Carey uses the label 'parallel individuation system' to describe what I am calling the Object-File System. Both these labels, as well as 'Object-Tracking System', are common throughout the numerical cognition literature. Importantly, all three labels refer to the same system discussed in chapter 1. Thus, Carey explicitly treats these as equivalent, writing "Parallel individuation is indeed the system of representation studied by Pylyshyn and others in their work on attentional indices… and by Kahneman, Treisman, Luck and many others in their work on object-based attention and visual short-term memory" (LeCorre & Carey 2008, 656). Similarly, Shea (2011) writes that "The starting point for Carey's developmental transition is the parallel individuation system ('object files' for short, although the system can also individuate events)" (Shea 2011, 6). Carey's use of the label 'Analog Magnitude System' does not reflect any important discrepancy with my use of the label 'Approximate Number System', nor does it express a commitment to the ANS skepticism described at the end of chapter 2.

[108] In some cases, Carey seems to use the implicit/explicit distinction to refer to whether or not a representation is accessible to consciousness. In other cases, a neural symbol or representation is explicit in that it is available as input for more central processing. For example, we can see that availability for further processing is key to Carey's use of the term 'explicit' in Carey's description of the ANS in the following excerpt:

> The evidence points to a system of representation in which number is encoded in the brain by some neural quantity that is a linear or logarithmic function of number. If this is right, we can say more about what is represented explicitly and what implicitly by this system. The symbols themselves are explicit. They are the output of the input analyzers and are available to central processors for a wide variety of computations. They can be bound to sets of quite different types of individuals. And various arithmetical computations are defined over them—numerical comparisons, addition, subtraction, and ratio computations. But much of the numerical content of this system of representation is implicit. There is no explicit representation of the axioms of arithmetic, no representation that 1–1 correspondence guarantees numerical equivalence. These principles are implicit in the operation of the input analyzers and in the computations defined over analog magnitudes, but they need not be available for the child to base any decisions on. (Carey 2009,135)

For more on the implicit/explicit distinction, see Ball 2016.

'a' selects one member from a collection, 'some' selects at least one member from a collection, and 'all' selects every member of that collection. The content of these quantifiers is thus radically different from what is produced by the OFS and the ANS.

First, consider the OFS. As I mentioned in section 1.4, while the OFS can support computations such as SAME/DIFFERENT, this is not the same as producing SOME/ALL/NONE, which, Carey claims, has explicitly quantitative content. Also, the OFS can produce symbols for specific individuals and track them, and distinguish between collections based on their numbers of individuals, but its range is severely limited. Of course, quantifiers like SOME or MANY clearly apply beyond the subitizing range. Further, the OFS does not produce content that explicitly distinguishes between one object file and many object files. Rather, its job is only to produce object files, not to quantify over them. What is more, the singular-plural distinction explicitly deals with the distinction between a singular individual and any other collection of objects, which suggests that it recruits different quantification processes from those that can be used over the representations of the OFS, which does not treat any collection containing more than one object as equivalent. Carey sums it up thusly: "the system of parallel individuation has no symbols for quantifiers, not even *one* versus *some*, and it has an upper limit of sets of 3" (Carey 2009, 256).

As for the ANS, while this system does produce explicit representations with quantitative content beyond the subitizing range, this content does not seem to apply to the same domain as natural language quantifiers. Rather, the ANS could be described as producing representations with content like THIS MANY, or THAT MANY. But this sort of content does not explicitly distinguish between objects that have a certain property and those that do not, like 'some' or 'most' would, for example. Linguistic quantifiers draw boundaries in the numerical domain. For example, the singular-plural distinction draws a line between one and any number larger than one, treating any collection containing more than one object as equivalent instances of plurality. The

ANS does not treat collections with different numbers of objects as equivalent. On the contrary, one of its main jobs is to detect differences in quantities of objects. The representations produced by the ANS make information about the approximate number of items in an attended set available to more general processes, but the grouping of items that allowed this explicit representation is not itself available for more general processing. Thus, the ANS "fails to have to the representational force of set-based quantification" (Carey 2009, 255).

Perhaps the OFS and ANS could team up to allow the quantificational distinctions expressed in natural language quantifiers? There is evidence against this possibility, since the OFS and the ANS are not well integrated, at least at first, as demonstrated by animal and infant failure to distinguish between quantities that straddle the subitizing range. In particular, failure of infants to retrieve 4 vs 1 crackers (Feigenson & Carey, 2005) is striking evidence that a singular-plural distinction would be used if it were possible to do so using the resources of the ANS and/or OFS. The fact that subjects fail at this task suggests that neither the OFS nor the ANS can allow infants as old as 20 months to grasp the difference between 'one' and 'more than one'. Carey sums up the situation thus:

> These data are consistent with the possibility that not only are there no explicit symbols for plural in the two core cognition systems with numerical content discussed in chapter 4, neither are there computations that treat all sets greater than 1 as equivalent and different from one. (Carey 2009, 256)

So, given Carey's evidence that the ANS and the OFS cannot account for the quantitative distinctions found in natural languages, a distinct cognitive system is needed to account for these. Carey calls this the *set-based quantification system.* It produces "explicit symbols with the content set and individual, plus distributive and collective computations over those symbols, to capture the meanings of natural language quantifiers, including even the singular/plural distinction" (Carey 2009, 256).

In English, the first set-based quantificational distinction children learn is the singular-plural distinction, as in "are X" vs "is X", or 'are xs' vs 'is an x'(Brown 1973). English-learning infants come to understand the distinction between one and many in searching tasks at roughly the same age as they understand the linguistic expression of this difference, so the question may be asked whether language is the source of this understanding. A good way to test this is to compare languages that mark plurals to varying degrees, since if language were the reason why children learn to distinguish between one and many, then children learning languages that do not emphasize this difference as much as others could be expected to develop this understanding later in life. Experiments comparing children's ability to distinguish one and many objects in searching tasks have found that infants begin successfully performing these tasks at roughly the same age, even in classifier languages like Chinese and Japanese that lack singular determiners and do not mark the plural of verbs and nouns (Li et al. 2009). According to Li et al., "experiments suggest that knowledge of singular–plural morphology is not necessary for deploying the nonlinguistic distinction between singular and plural sets" (Li et al. 2009, 1644). Given that behavior manifesting the singular-plural distinction appears to be independent of the language one learns and appears to show up at around the same age, regardless of language, Carey concludes that "set-based quantification is part of the machinery children bring to the task of language learning, either as part of the language-acquisition device or as part of general representational capacities" (Carey 2009, 261).

Evidence from 15-month old infants (Barner et al. 2007) and rhesus macaques (Barner et al. 2008) also suggests that, under the right circumstances, these subjects can distinguish between collections containing single individuals and groups of up to 5 individuals in food choice and simple habituation tasks. In such tasks, Carey interprets the reaction as based on a distinction between one and many, since they succeed at tasks including 1 vs 2, 1 vs 3, 1 vs 4, and 1 vs 5, but fail at 2 vs 3, 2 vs 4, and 2 vs 5. In short, "there is no evidence that learning explicit linguistic representations for set-based quantification underlies the changes observed in English learners" (Carey 2009, 260). Given that this behavior is present in preverbal infants and animals, this counts as further strong evidence that language is not required for the acquisition of set-based quantification. Also, none of the quantity representations of the ANS or the OFS are behind this behavior, given that subjects would have succeeded at some of these ratios if one of these systems had been in charge.

I have just reviewed Carey's evidence for the existence of a third representational system that supports quantity-related processing, which Carey calls set-based quantification. While this module may underlie natural language quantification, its presence in animals and preverbal infants makes it likely that it is an innate language-independent representational resource with conceptual content. Unfortunately, the mere fact that systems like set-based quantification, the ANS, or the OFS can be recruited to process information about quantities of discrete objects in our environment does not allow us to understand how these systems are responsible for our more arithmetical abilities. And yet, we also saw in chapters 1 and 2 that both the ANS and the OFS do appear to be recruited in numerical cognition, as evidenced by distance effects and limits to the subitizing range, to name a few places where their signatures are visible in behavioral data.

This continued presence throughout our lifespan would suggest that these systems must somehow be responsible for our numerical abilities, but that they undergo a profound transformation in the ontogeny of numerical content. I presented some of Carey's reasons for doubting that a mapping of the sort proposed by Dehaene can explain this transformation in the previous chapter. What is missing is an account of how attributing the core cognition status to the ANS and the OFS as well as to set-based quantification allows Carey to succeed where others have failed, in explaining how we develop representations with content that is discontinuous, in important ways, with that of its building blocks.

Now that we are familiar with the properties of core cognition and how these could explain how innate perceptual analyzers could produce content that figures into rich, domain general processes, we are in a position to take a closer look at how Carey wishes to bridge the gap between the content of core cognition modules and fully-fledged concepts of number.

In so doing, we will be looking to see how Carey both describes actual cases of discontinuities and whether she provides a plausible explanation of how such discontinuities arise in development. The next section summarizes Carey's argument for the existence of such discontinuities, while section 4.7 presents her account of learning natural numbers from core cognition, which she presents as evidence of the important role of core cognition in explaining how discontinuities can be overcome.

4.6 Learning through discontinuity?

Many arguments have been proposed by philosophers through the ages that endorse what Carey calls the *continuity thesis*, according to which "all the representational structures and inferential capacities that underlie adult belief systems either are present throughout development or arise through processes such as maturation" (Carey 2009, 18). One well-known vocal proponent of the continuity thesis is Jerry Fodor (1983; 1997), whose views on the matter are well summarized by Carey: "one cannot learn what one cannot represent" (Carey 2009, 18). Fodor has argued that despite superficial differences, each form of learning can be described as an instance of hypothesis testing. Given that it is impossible to test a hypothesis without first having a representation of its content, Fodor concludes that it is impossible to learn new representations, since any learned representation was already in our possession in the form of a yet-to-be-proved hypothesis. This leads Fodor (1998) to endorse a radical concept nativism (and atomism).

To help frame the plausibility of the continuity thesis, Carey cashes out talk of representational systems being discontinuous with one another in terms of their respective representational or expressive power: when one system has incommensurably more expressive power than the other, it is considered as discontinuous with it. One way to show discontinuity between two systems is to show that one is "qualitatively more powerful than" (Carey 2009, 288) the other. For example, if one system is unable to describe phenomena that can be described in the vocabulary of the other, then the more expressive system is discontinuous with the other. Examples of such discontinuities include representations of North in indigo buntings and those mentioned in chapter 1 in the discussion of object permanence (section 1.3.3).

Carey argues that the continuity thesis is wrong, since it fails to explain how learned concepts could ever acquire content that is incommensurate with that of its building blocks. And yet, there are countless examples of such learning through discontinuities. For example, as Carey points out, Piaget (1980) already challenged Fodor's denial of conceptual development that transcends discontinuity when he argued that advanced mathematical notions such as complex or imaginary numbers cannot plausibly be considered the outcome of natural maturational development, given that most mature adults do not entertain such concepts. Going further than Piaget, Carey claims we need not look any further than the natural numbers to prove Fodor wrong, since the natural numbers, which she (and many others) considers a cultural construction, have more expressive power than any of the systems used to construct them.

One of the major theses of Carey's 2009 monograph is that it is possible for humans to develop representations that are discontinuous with the representations they are built from. Carey does not disagree that learning can usually be described as a form of hypothesis formation, nor does she deny that learning builds new representations out of old ones. For Carey, the main problem with Fodor's view is that it is false: the literature on cognitive development, which Carey deploys masterfully throughout her 2009 book, is filled with cases of children acquiring novel representational resources that she claims are discontinuous with their building blocks, in the sense that it is impossible to express the content of the learned concept using the content of the concepts used to learn it.

According to Carey, the continuity thesis cannot apply to the relation between the representations of core cognition and those of natural numbers because no system of core cognition represents natural numbers. As Shea puts it,

> On one side of the conceptual discontinuity are various core cognition systems with number-related content. On the other side is a system of representations of natural number. (Shea 2011, 6)

The very fact that we are trying to explain how we can build an understanding of natural numbers from the content of core cognition seems to imply that it is possible to learn conceptual content that is discontinuous with its building blocks, given that the core cognition systems that appear to have numerical content, either explicitly (ANS) or implicitly (OFS), cannot produce content with which we could ever represent anything as precise as 27, or even 6, no matter how many instantiations of these we are exposed to. Worse, these systems cannot account for a successor function outside the subitizing range. Similar reasons prevent natural language quantifiers from fitting the bill, given that these are, in some sense, more coarse even than the representations of the ANS, given that they can only distinguish between singular and plural, and sometimes include symbols for pairs and triads, which suggests both limited range and precision.[109]

---

[109] Of course, as Carey points out, "Numerals, when used in sentences, are quantifiers, so this system is clearly implicated in number word meanings" (Carey 2009, 296), but these symbols are not available innately, which means we cannot appeal to them as being the source of the representation of natural numbers.

The gap problem, then, is about discontinuity between the content of representational systems and how we develop novel representational resources. In this sense it is more complicated than merely mapping a word onto a pre-existing representation:

> Unlike learning the meaning of the plural morpheme, which is a problem of finding a mapping between antecedently available nonverbal representation and a linguistic expression, learning the meaning of "seven" requires constructing a new representational resource. (Carey 2009, 305)

This explains why Carey thinks developing representations of natural numbers should be harder than merely learning the meaning of plural morphemes, for example, since in the case of plurals, we are born with the representation to which the linguistic item must be mapped, as seen in the previous section. The fact that Wynn's (1992b) data on children learning the meaning of number words discussed at the end of chapter 3 involves such a lengthy, protracted learning process, as well as the fact that children are subset-knowers for such a long time, is taken by Carey as evidence that learning what natural numbers are involves building representational resources discontinuous with its building blocks.

An important implication of the continuity thesis for conceptual development is that it is impossible to develop representations whose content cannot be described or articulated with the same terms used to describe the content of the representations out of which it is built. If this is true, then it should be possible to frame the development of our understanding of natural numbers in terms of maturational effects on innate cognitive systems. Given the limitations of the core cognition systems that appear recruited in numerical tasks, there is little reason to think this can be done.

And yet, as seen in the previous chapter and in chapter 2, the fact remains that external representations of numbers like Indo-Arabic numerals and number words do eventually get mapped to the ANS, though Carey has provided convincing evidence that this mapping occurs after the crucial gap-bridging learning mechanism has taken place.

This begs the question of how we managed to bridge the gap before mapping external symbols for numbers onto the ANS. For while it is one thing to show that there exist conceptual discontinuities by interpreting developmental data, it is another thing to explain how such discontinuities occur. If we are to properly explain how we bridge the gap between core cognition and theories of numerical content, we must have a learning device capable of explaining how such discontinuities can arise. This is what Carey dubs *Quinian bootstrapping*.

4.7 Quinian Bootstrapping

4.7.1 Meaningless lists of symbols

In the previous section, I sketched a few reasons explaining why Carey interprets the limitations of the output of our core cognition modules as evidence that they are discontinuous with the content of representations of natural numbers. To explain how core concepts allow us to develop fully fledged theories or learn representations like NATURAL NUMBER, it is important to understand how some representations can be built out of others in a way such that the more complex representations have content that cannot be expressed or understood using representations produced by its building blocks.

To describe such a learning trajectory, Carey claims we need to characterize the initial stock of representations as well as the learning processes that modify these. In the case of learning number concepts, for Carey, the initial state is a motley crew of representational systems of core cognition: the OFS, set-based quantification, and the ANS, "the three systems of representation with numerical content that are bequeathed to human beings by natural selection" (Carey 2009, 304-5).

As for the learning mechanism, Carey's account needs a learning mechanism capable of building representations whose content is incommensurate with its building blocks. This is what Carey calls *Quinian Bootstrapping* [sic],[110] which Shea describes as "the process whereby a set of uninterpreted symbols interrelated by a network of inferential dispositions are connected up to the world so as to acquire a meaning" (Shea 2011, 6).

For Carey, the problem of how children acquire representations of natural numbers can be decomposed into three sub-problems. The first involves how children learn the list itself. The second is how they learn the meaning of each symbol in the list, and the third is how they learn that "the list itself represents number, such that the child can infer the meaning of a newly mastered numeral symbol (e.g., "eleven") from its position in the numeral list" (Carey 2009, 308).

The first step in Quinian bootstrapping is the learning of lists of symbols which, at first, are meaningless:

> In Quinian bootstrapping episodes, mental symbols are established that correspond to newly coined or newly learned explicit symbols. These are initially placeholders, getting whatever meaning they have from their interrelations with other explicit symbols. (Carey 2011, 120)

In the case of numbers, this is the count list (1, 2, 3, 4…), which children learn to recite long before they have any developed number concepts. As seen at the end of chapter 3, it is only many months after they have learned this list that children start correctly applying the word 'one'. Months later, they form the ability to use 'two' and, after another while, 'three' (Wynn 1990).

There is ample evidence that children can learn ordered lists of symbols without

---

[104]Quinian bootstrapping shouldn't be confused with the syntactic and semantic bootstraps often discussed in language learning literature, since only Quinian bootstrapping involves discontinuity and the construction of novel representational resources, as opposed to linguistic ones. This is not true of semantic and syntactic bootsrapping, where infants exploit linguistic categories and distinctions to figure out the meaning of words using their semantic and syntactic environments.

grasping their meaning – assuming there is one. Obvious cases are the alphabet and segments of numeral lists, as well as days of the week and months, but other examples include nursery rhymes like 'eeny-meeny-miny-moe', which, apart from being sounds placed in a certain order, do not seem to have any particular meaning attached to them.[111] This illustrates the first stage of Quinian bootstrapping, during which the only meaning attached to symbols are their interrelations – in cases like learning natural numbers, the order in which the sounds must come. The bootstrapping metaphor emphasizes the fact that when learning is due to Quinian Bootstrapping, the structure we are building is initially not grounded. At first, all we have is a skeleton that shows the relations between the symbols. Once we have the frame of the structure built by the relations between its symbols, then we can add semantic meat to it.

It is in the second stage of Quinian bootstrapping, where the meaningless numeral list is gradually interpreted by combining representations of core cognition, that Carey's account differs from Dehaene's and others who take the ANS to be the only system recruited in the learning process. We saw that data concerning the lengthy and gradual process during which children become cardinal-principle knowers imposes a descriptive constraint that any bootstrap – or other account of the ontogeny of natural number concepts – must satisfy, that of explaining how partial meanings for numeral lists are constructed out of previously existing representational content in a way that accounts for the piecemeal, gradual process of going through stages of subset-knowing. We also saw that accounts like Dehaene that try to explain how we bridge the gap by proposing a mapping to the ANS do not fare well when faced with this constraint. These details of subset knowers are crucial empirical constraints that guide Carey's Bootstrapping account of how we bridge the gap. To avoid the pitfalls associated with mapping to the ANS, on Carey's account, interpreting numeral lists does not recruit the

---

[111] Shea (2011) observes that the 'elemeno' segment of the alphabet is often interpreted as a single item by children learning the alphabet.

ANS. Instead, “the resources of parallel individuation, together with those of set-based quantification, underlie the construction of the numeral list representation of number” (Carey 2009, 309).

Wynn's developmental data raise many questions: for each stage of subset-knowing, where does the partial knowledge come from? Why does it take such a long time for children to generalize this knowledge to numerals larger than four? Also, why does the transition from subset-knower to CP-knower regularly take place after children have become three-knowers (or, in rare cases, four-knowers)?

The fact that the transition from subset-knower to CP-knower occurs after children have become three- (or four-) knowers marks a rather elephant-sized hint that parallel individuation may be involved in this process. And yet, as seen in chapter 1, parallel individuation has no explicit symbols for numbers. While Wynn’s developmental data have parallel individuation’s fingerprints all over them, there is no obvious route from representing individuals to providing numerical content for numerals. How do we explain this? The answer, according to Carey, is in set-based quantification and how linguistic cues raise the saliency of its representations. Since this system has quantitative content in the subitizing range, perhaps, combined with parallel individuation, it could supply the meaning for the various stages of subset-knowing. I outline how this could happen, and the evidence supporting this hypothesis, in the next section.

4.7.2 Partial meaning and quantifiers

Carey claims there is strong evidence that natural language quantifiers and set-based quantification play a major part in giving numerical meaning to an initial segment of the numerals in the counting routine memorized by the child. Cross-linguistic studies can be used to test this possibility: if natural language quantifiers can ground the meaning of numerals in the representations of set-based quantification, then languages with more quantificational markers could facilitate this semantic association by making the distinction between 'one' and other numerals more salient, for example. This does appear to be the case, since classifier languages like Japanese and Chinese lack common singular-plural markings, and so do not highlight this quantitative distinction with the same force as does English (nor Russian). This means that if children exploit this distinction to bootstrap the meaning of the first few numerals, children learning languages that do not highlight it as much should take more time to learn the meaning of the first few numerals. This is exactly what was found: while children growing up in Russian, Japanese, and English-speaking communities all learn to recite an initially-uninterpreted numeral list by the same age, researchers found that Japanese children were slower to become one-knowers (Sarnecka et al. 2007). This is likely due to the absence of morphemes with content associated with the singular-plural distinction. Similar results were found in Mandarin-learning children (Li et al. 2003), which suggests that the absence of a morphological marking for quantity distinctions retards the onset of subset-knowing.

According to Carey, the hypothesis that numerals are first interpreted as quantifiers is further supported by historical linguistic data according to which the first few number words have special status across languages throughout history (Hurford 1987). The first few numbers are more commonly used and more likely to be involved in noun-phrase syntax, as illustrated by the fact that in many languages, dual markers (e.g. 'bi-', 'both', 'pair', in English) are less common than single markers, while triple markers (e.g."tri') are even less common, but still more common than markers for larger quantities.

This story of how we come to learn the meaning of the first few words in our count list due to the saliency-increasing effects of natural language quantifiers on the representations of set-based quantification seems well supported. If the first hypotheses that subset-knowers bring to the problem of figuring out what number words in their count list means comes from the resources of set-based quantification, then languages that emphasize the distinctions of set-based quantification should facilitate confirming this hypothesis.

4.7.3 Enriched Parallel Individuation

However, assuming this story works so far, we still have work to do to explain how we bridge the gap, since this explanation of partial meanings does not tell us why the generalization of CP-knowers occurs at the boundary of the subitizing range. Also, what sort of numerical content can we get from plural markers, and why would quantifier-based learning display the limits of parallel individuation? To answer these questions, Le Corre and Carey (2007) propose enriched parallel individuation (EPI).

To understand EPI, we must first remember that the OFS supports one-to-one correspondence (11C hereafter) between states of object tracking, which can be described in terms of how many object files are open in a given task. Despite the fact that the OFS is dedicated to tracking objects, Carey claims that this system nevertheless supports comparison between collections of individuals, as evidenced by the many studies mentioned in chapter 1, which show that infants and preverbal animals can choose between two scenarios based on the number of objects they contain, even when the total number of objects in both scenarios exceeds the set-size limit of the OFS:

> evidence that infants can use parallel individuation to choose a set of three crackers over a set of two crackers (Feigenson et al., 2002) implies that the parallel individuation system somehow keeps track of which set is cracker, cracker, cracker and which is cracker, cracker. (Le Corre & Carey 2007, 657)[112]

---

[112] Similarly, Le Corre & Carey claim that "preverbal infants can make models of at least two sets of

Assuming that the OFS can support such computations that compare collections of objects, a record of the states of the OFS used in such comparisons must be kept in memory in order for comparisons between states to take place (Shea 2011). However, the associations between words and states of the OFS need to be stored in long-term memory because the representations of the OFS are generally temporary, given their tracking purposes (Le Corre & Carey 2007). The enriched part of this parallel individuation thus refers to the fact that object files are boosted with a long-term memory model of a state of the OFS.

This long-term memory record is important because it can be used to build an association between the collection of objects being tracked and the quantifier that applies to that collection:

> The child makes a working-memory model of a particular set he or she wants to quantify (e.g., {cookie cookie}). He she then searches the models in long-term memory to find that which can be put in 1–1 correspondence with this working-memory model, retrieving the quantifier that has been mapped to that model. (Carey 2009, 324)

The discussion above of the role of set-based quantification in grounding a partial understanding of the list of number words showed that words can become associated with representations of set-based quantification due to the saliency-increasing effects of repeated exposure to these words while tracking the relevant number of objects. This means that representations for single individuals (e.g. {OBJECT}) can be associated with linguistic symbols for individuals (e.g. 'one', 'a') due to repeated exposure to such external symbols when tokening single object files. Similarly, symbols for two

---

individuals, each subject to the set-size limit on parallel individuation, and hold both in working memory at once" (2007, 656). According to Gallistel (2007) it is unclear how the OFS can support such a computation, since its role is limited to constructing representations of individuals, not sets of these. See Shea 2011 and Beck 2016 for some clarifications of this aspect of EPI. For present purposes, such details need not concern us, given that our main problem with Carey's account lies in more general matters, discussed below in section 4.8.

individuals (e.g.{OBJECT}, {OBJECT}) could be associated with linguistic symbols for pairs, and the same sort of association can be applied to the word 'three' and the state in which the OFS is when it track three objects. This way, the numerical content of set-based quantification gets associated with a linguistic label for this content, and stored in a long-term memory model of the state of the OFS during this association.

We saw above that English-speaking children first learn the singular-plural distinction, which helps them get the content for 'one'. The fact that both 'one' and 'a' map onto {OBJECT} in long term memory explains how 'one' gets its meaning: set-based quantification supplies the numerical meaning of 'a' and 'one' because these terms are most often used when the child opens a single object file. This means that the association between the numerical content of set-based quantification – here, the distinction between singular and plural – and the number word gets stored in an object file representation for OBJECT. Then, children notice that words like 'some' and number words other than 'one' are applied to collections that contain more than one object, which allows them to treat number words like quantifiers. This can explain why one-knowers will respond to 'how many?' questions using number words instead of, say, color words, even if they use number words outside their subset-knowledge range randomly: number words make up the initially-meaningless skeleton that now has one grounded member, and the child knows that only a word from this list will do. This step is followed by one where children learn that 'two' only applies to a subset of pluralities, namely, those that contain two individual files. The word 'two' then gets mapped to a long-term memory model of pairs of objects, and the same process allows children to map the meaning of 'three' to a long-term memory model of three individuals.

The general idea here, then, is that children rely on "models of the sets of individuals held in parallel in working memory," (Carey 2009, 327) where each model can be thought of like an abstract description of what we do when individuating of up to four objects in parallel. According to Carey, "All of the computational resources required for enriched parallel individuation are known to be available to prelinguistic infants" (Carey 2009, 324). This includes the ability to put two different collections of objects into 11C, treating collections as objects, and quantifying over collections with natural language quantifiers.

4.7.4 The crucial analogy and CP-knowledge

Let us assume Carey's explanation works up to this stage. If so, it looks like we have an idea of how content gets attributed to the first few numerals. Children who have mapped the meaning of the first few number words are then poised to notice analogies about numeral order and numerical content within the subitizing range, since they have learned to label states of the OFS based on the number of object files being used to track objects. They are then in a position to notice that when counting two or three objects, the last word used is also the word that describes how many objects were counted. At this point, 'all' that is left is for children to do is generalize this counting principle to the number words to which they have yet to assign any meaning.

Here, the child must notice an analogy between number word order and order in long-term memory models for individual files: taking one step in the numeral list is analogous to "next in the series of models ({i}, {j k}, {m n o}, {w x y z}) related by adding an individual" (Carey 2009, 327), or "next state after additional individual has been added to a set" (Carey 2009, 328). While this analogy only applies to numerals 1-4, it still can serve as the basis for the crucial induction, where the child realizes that " if 'x' is followed by 'y' in the counting sequence, adding an individual to a set with cardinal value x results in a set with cardinal value y" (Carey 2009, 327). This critical analogy, then, allows the child to understand the successor function, with which the meaning of any number word in the count list can be interpreted.

There are, then, two essential elements to Carey's account of learning via Quininan Boostrapping: first, the elements of the system to be learned are initially placeholders whose meaning is, at best, partially interpreted by previously available resources. Second, as for how these symbols become fully interpreted, Carey takes it that the process responsible for the full interpretation of symbols is one of noticing an analogy and coming to a crucial induction that allows the child to generalize the hard-won associations built during the construction of the partial mappings of subset-knowing and understand that going one step further in a count list means adding one object to the collection being counted.

4.8 The origins problem

As Jacob Beck points out, "Carey's account of Quinean bootstrapping has been heavily criticized" (Beck 2016, 110). For example, critics have claimed that Carey's bootstrapping-based account is circular (Fodor 2010; Rey 2014) and that it underspecifies the potential conclusions that children could come to when making the crucial induction (Rips et al. 2008a; Rey 2014). Others have complained that it is unclear just what episodes of Quinian bootstrapping are and that the notion of enriched parallel individuation is equally confused (Gallistel 2007). Worse, there is evidence that some children manage to learn the meaning of 'one' despite applying 'a' to collections containing more than one object, which suggests that the presence of singular-plural markers may not facilitate acquisition of the meaning of 'one' (Barner et al. 2009), contrary to what Carey claims. While these theory-specific issues may turn out to be genuine worries for Carey,[113] they will not concern us here.

---

[113] While Shea (2011) clarifies certain important aspects of Quinian Bootstrapping, Beck (2016) attempts to clarify what bootstrapping is in order to counter the circularity and ambiguity objections.

The problem I have with Carey's account is one that can also be levied against any account that, like Carey's, attributes an essential role in external symbols for numbers in explaining how we bridge the gap. In Carey's bootstrapping account, lists of number words play two crucial roles in explaining how we develop an understanding of natural numbers. Carey's solution to the gap problem only goes through by relying on the availability of a numeral list, both as the initial placeholder and in the saliency-increasing effects of structured and ordered number words in associating states of the OFS with number words: "At the heart of Carey's theory is the idea of memorizing a set of relations amongst uninterpreted symbols" (Shea 2011, 7). Without these lists of number words, the learning process could not go through: "Quinian bootstrapping processes require explicit symbols, such as those in written and spoken language or mathematical notational systems" (Carey 2009, 306).

The problem I want to talk about is this: how can we rely on external symbols for numbers in our explanation of the development of numerical content when the existence of such symbols in turn depends on the existence of numerical content?

For starters, consider the role played by integer lists in Carey's account. The problem here is, how could such lists possibly emerge without someone first having had some kind of number concepts? The fact that someone (or, much more likely, many people, in a gradual process of personal innovation and cultural transmission) had to come up with this counting routine means that it was possible to think about numbers (or perhaps, more basically, about precise quantities) without relying on lists of words. This in turn suggests that there can be individual-level development of number concepts without such developed external support, which would seem to imply that some other artifact-free learning processes are involved in the origin of basic number concepts. But if we must appeal to such aretefact-free processes to describe the origin of number concepts, then how can we consider Carey's bootstrap to be an explanation of the origins of natural number representations, given that this account depends on external crutches whose absence did not prevent people from developing an understanding of natural numbers in the past?

My claim is that, if we are trying to explain the ontogeny of number concepts, our theory should apply to *everyone* capable of thinking about numbers. But since some people seem to have been able to think about numbers without external aids in the (distant) past, any account that depends on such support will not apply to every case of numerical cognition. At best, such externalist accounts could describe how numerical cognition emerges in a numeral-enriched environment.

Importantly, while this is a problem for Carey, it generalizes to any explanation that relies on the presence in the learner's environment of an organized system of symbols or practices for precise quantification. Thus, Dehaene's mapping, which relies on number words and symbols to explain how we develop an understanding of natural numbers, is also a target here.

This may seem unfair to Carey: after all, her account is meant to explain the ontogeny of natural number representations *in a numeral-enriched world*:

> Like virtually all researchers in this field, we agree with Gallistel and Gelman that the verbal numeral list deployed in a count routine is the first explicit representation of the natural numbers mastered by *children growing up in numerate societies*. (Le Corre & Carey 2008, 651)

However, the fact that it is possible to develop some basic number concepts without the type of external support used in Carey's model seems to suggest that cases that do involve external support might somehow appeal to a more fundamental process, which Carey's externalist framework is leaving out. So while it may seem unfair to Carey to criticize her for not taking into consideration historical development, given that her theory is aimed at the individual, ontogenetic level, there is reason to do so: the ontogeny of number concepts in a world where symbols for numbers abound cannot be completely separated from past cases of numeral-free ontogeny, since the former depends on the latter in important ways.

This line of thinking was perhaps behind Overmann et al.'s (2011) comment on Carey's *précis* of her 2009 book, where they ask, "In the absence of a numeral list, how could a concept of natural number ever have arisen in the first place?" (Overmann et al., 2011, 142) Carey appeals to extended cognition in a cultural setting to answer this question:

> Understanding the invention of tally systems would involve understanding how people came to the insight that beads could serve this symbolic function, rather than decorative uses, or as markers of wealth, or myriad others. That is, the availability of an artifact that could serve as the medium of a tally system doesn't explain how it came to be one. Now that we are in the realm of speculation, I believe, contra Overmann et al., that body counting systems could well also play an extended cognition role in the cultural construction of integer representations. (Carey, 2011, 159)

Unfortunately, it is difficult to see how this can tell us the whole story, since this simply raises the question of how these body-counting systems emerge, if not by some kind of purely blood-n-bones intracranial learning process. Assuming that our first number symbols were parts of the body, how did we come to point to these and intend to communicate numerical content without first having developed this numerical content by some other cognitive process that did not rely on external symbols for number? For example, how could a person point to their right knee or left thumb to communicate 'six' without first having come to some understanding of SIX? Whatever process allowed such symbols to emerge, shouldn't *that* be where we look to find the key to the mystery of the development of numerical content? Carey offers no reason why the crucial analogy and accompanying induction could not occur in different, numeral-impoverished contexts.[114]

---

[114] While Coolidge & Overmann (2012), Overmann (2015), and Malafouris (2010) offer similar lines of reasoning concerning the emergence of numerical cognition in numeral- impoverished environments, their criticism is mostly aimed at the role of language in this process. Importantly, their proposed accounts also rely on external symbols—either in the form of fingers or clay tokens—to explain how we

The success of Carey's bootstrapping account of how we bridge the gap – as well as many other externalist accounts, as discussed in the next chapters – thus hinges on whether or not appealing to extended cognition and cultural evolution can explain how it would be possible to bridge the gap in a world where no count lists, integers, or experts willing to teach us mathematical practices abound. Exploring the role played by external objects and cultural evolution in the development of representations of natural numbers will occupy the last few chapters of this thesis. The first step is to explore and set up what extended cognition is and how it can apply to the development of numerical cognition. I do this in the next chapter.

transcended the limitation of our innate cognitive systems. The point I am making here goes further and applies to any external representation with numerical content, including numerals and number words, but also clay tokens, body parts, and other 'enactive signs' (Malafouris, 2010). I explore this distinction in more detail in the next chapter

CHAPTER V

EXTERNALISM AND NUMERICAL COGNITION

5.1 Introduction: external supports for numerical cognition

In the past two chapters, we have seen two attempts to bridge the gap between our representations of natural numbers and the content produced by evolutionarily-inherited systems like the ANS, the OFS, and, in Carey's case, set-based quantification. We saw that Dehaene's triple-code model provided a plausible description of the way in which numeration systems like the Indo-Arabic numerals and linguistic symbols for numbers come to invade different parts of our brain, but that his description of how we learn the meaning of these symbols by a mapping to the representations of the ANS does not seem to account for Wynn's (1990, 1992b) and Le Corre & Carey' (2007) developmental data. We also saw how Carey's account of this learning, whose focus is less on the neuronal details and more on the developmental data, offers a plausible account of how we can develop representational systems discontinuous with their building blocks by an induction over the meaning of the first few number words, which is made possible by representations of set-based quantity and object files.

While Dehaene and Carey disagree on many points – especially the role of the ANS in the development of our understanding of what natural numbers are – both their accounts heavily rely on the presence of developed numeration systems to bridge the gap. Despite their differences, then, both can be considered externalists about numerical cognition, in the sense that their accounts suggest that external symbols and artefacts are ineliminable constituents of our ability to think about – and with – numbers. This view, according to which mathematics depends on external objects, is not new. For example, there is reason to believe that Kant held that writing was constitutive of mathematical practice (Macbeth 2013). However, it is important to be clear about just what sort of dependence we are talking about here. After all, most people would agree that oxygen is also necessary to practice arithmetic, since mathematicians need to breathe, but few would argue that oxygen is a constituent component of mathematical practice. Thus, we need to know how externalist accounts of numerical cognition frame the dependence of arithmetical practices on external symbols and artefacts in order to determine to which extent it is legitimate – necessary, even – to adopt an externalist perspective on numerical cognition.

As I mentioned at the end of chapter 4, externalist accounts of numerical cognition share a problem concerning the origins of our numerical abilities, due to their reliance on external symbols and artefacts: since it would appear reasonable to think that numeration systems are human constructions, it looks like they themselves require an ability to think about numbers in order to come into being, which means that it should be possible to think about numbers without external symbols, on pain of regress of external supports. There are, of course, many ways to counter this potential problem. Carey's response to Overmann et al's (2011) comment, where she claimed that "counting systems could well also play an extended cognition role in the cultural construction of integer representations" (Carey 2011, 159), displays two important externalist options. First, the externalist can deny that the origins of numerical artefacts and symbols is in purely-internal representations by adopting an extended approach to cognition. This is the focus of the present chapter. Or, alternatively, the externalist could appeal to cultural evolution to explain the historical development of numerical practices, and thus negate the importance of the origins problem. This possibility is explored in chapter 6.

In section 2 of this chapter, I discuss the extended approach to cognition, as introduced in Andy Clark and David Chalmer's (1998) article, 'The Extended Mind', followed by Clark's (1998) views on explanation in extended cognitive systems, which I will use to evaluate externalist approaches to the gap problem. I then move on to externalist approaches to numerical cognition in section 5.3, focusing on Catarina Dutilh Novaes' (2013) treatment of constitutivity in extended systems of numerical cognition in section 5.4, before posing a few challenges to this approach in section 5.5. To help ground this discussion, I take a look the historical origins of symbols for precise quantities in Sumeria in section 5.6, while section 5.7 offers a perspective on the limited numerical abilities of cultures like the Pirahã (Frank et al. 2008) and the Mundurucu (Pica et al. 2004). The main claim I make in this chapter is that externalist accounts fail to explain what makes the difference between an individual that has bridged the gap and one who has not in situations where both have access to the same external support for cognition, since the difference in these cases lies in our head.

5.2 'The Extended Mind'

5.2.1 Thinking outside the box

David Chalmers, in his introduction to Andy Clark's (2008) monograph, *Supersizing the Mind*, writes this about the idea that the mind can extend beyond the brain: "The thesis has a long history: I am told that there are hints of it in Dewey, Heidegger, and Wittgenstein. But no one has done as much to give life to the idea as Andy Clark" (Chalmers 2008, x). This is certainly true. For while the notion of extending cognition beyond the brain had received attention prior to Clark's work on the topic (e.g. Hutchins 1995; Wilson 1994; Haugeland 1993), it was only following the publication of Clark and Chalmers' 1998 article that the possibility of extending the mind beyond the skull truly took off as a major topic in the philosophy of mind.[115]

In this highly influential article, Andy Clark and David Chalmers (ACDC hereafter) question the bounds of cognition, asking where the mind stops and the world begins.[116] Deploying compelling intuition pumps and thought experiments, they argue that, in some cases, parts of the mind lie outside our heads – hence, the title of their article, 'The Extended Mind' (TXM hereafter). Contrary to the traditional skull and bones localization of mind, TXM allows a motley crew of external objects to be considered constitutive elements of cognitive processes. These include cellphones, notebooks, and anything else that can reliably be used and accessed to process information. This externalist framework focuses on the dynamic interaction between individual brains and the objects in their environment to understand how objects outside our heads, including culturally accumulated artefacts, can supplement our cognitive abilities. The implication is that the mind can extend beyond the barriers of our skull to include parts of our environment.

In a nutshell, the main motivation behind the extended mind is the realization that parts of the world outside our head play an active role in some cognitive processes, and that if we are to allow parts of our brain to be parts of our minds because they play such a role, then we should allow things outside our heads to be parts of our mind as well,

[115] It is worth mentioning that such extension of the boundaries of the individual had already taken place in other domains. For example, Dawkins' (1983) extended phenotype and Oyama's (1985) Developmental Systems Theory both questioned the bounds of individuals in biology.

[116] As Pöyhönen (2014) points out, in the fallout following publication of their article, which includes the development of many varieties of externalism, the conversation has largely treated the extended mind and extended cognition as synonymous. However, given that the mind can include things that aren't parts of cognition, this is a problematic equivocation. This being said, while I will try to focus on extended cognition, this terminological issue does not have significant impact here, given that the focus is not on the nature of mind itself, but the degree to which external artefacts can be considered constitutive parts of numerical cognition. What minds are is neither here nor there with respect to the origins of numerical cognition.

when they play such active roles. According to ACDC, failure to include such extracranial elements in cognitive loops on the basis of their geographical properties is tantamount to cranial chauvinism.

ACDC use a few intuition pumps to support the legitimacy of their controversial thesis. One often-cited source of evidence that some cognition takes place outside the head is the classic video game *Tetris*, in which the player must figure out the best way to place four types of geometrically-distinct shapes in order to fill the bottom of a screen. To figure out the best position in which to place the shapes, the player can press a button and the shape will rotate on the screen, thus making its ideal positioning visually salient. Kirsh & Maglio (1994) call the physical rotation of an image on a screen to determine if it can fit into a slot an *Epistemic Action.* Such actions "alter the world so as to aid and augment cognitive processes such as recognition and search" (ACDC 1998, 8). These contrast with *Pragmatic Actions*, which alter the world because such an alteration is desirable for its own sake.

To illustrate the cognitive benefits of this use of visual cues, consider the fact that rotating the image on the screen and pressing the button takes about 300 milliseconds, while rotating it in our head takes about 1000 milliseconds. This shows that there are some cognitive processes – in this case, figuring out the best way to place a geometric shape – that heavily rely on parts of the environment because they are far more efficient ways of getting the information we need.

Given that the rotation of the image is the result of our pressing of a button, and that the change in visual stimulus that results from pressing the button occurs on a screen instead of in our brain, there is a sense in which parts of our environment are processing information for us. We are coupled to the video game, in that we are caught in a loop of continuous reciprocal causation,[117] where what we do influences the image on the

[117] Clark describes the notion of continuous reciprocal causation as applying to cases where "some

screen, which in turn influences what we do next:

> the human organism is linked with an external entity in a two-way interaction, creating a *coupled system* that can be seen as a cognitive system in its own right. All the components in the system play an active causal role, and they jointly govern behavior in the same sort of way that cognition usually does. If we remove the external component the system's behavioral competence will drop, just as it would if we removed part of its brain. Our thesis is that this sort of coupled process counts equally well as a cognitive process, whether or not it is wholly in the head. (ACDC 1998, 8, italics original)

ACDC claim that if some actions involve use of external media to augment cognitive processes, then, in certain cases, these actions, by including external media, extend the mind into the world:

> Epistemic action, we suggest, demands spread of epistemic credit. If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process. (ACDC 1998, 8)

Other common examples of objects to which we can get caught in loops of reciprocal causation include using pen and paper to run through calculations and physically re-arranging letters when playing Scrabble: in both these cases, we are manipulating external media to complete cognitive tasks instead of doing it in our heads. ACDC claim that "In a very real sense, the re-arrangement of tiles on the tray is not part of action; it is part of *thought*" (ACDC 1998, 10, emphasis original).

ACDC accept that in some cases, like experiences, the mental state can be individuated internally. Other parts of the mind, however, seem like fair game to their externalist

---

system S is both continuously affecting and simultaneously being affected by activity in some other system O...we often find processes of CRC that criss-cross brain, body, and local environment" (Clark 2008, 24).

approach to cognition. ACDC's main case is that of belief:

> we will argue that *beliefs* can be constituted partly by features of the environment, when those features play the right sort of role in driving cognitive processes. If so, the mind extends into the world. (ACDC 1998, 12, italics original)

Their well-known example to motivate the idea that beliefs can extend into the world is the thought experiment involving Otto and Inga. Inga is an example of "a normal case of belief embedded in memory" (ACDC 1998, 12). She hears there is an exhibit at New York's Museum of Modern Art (also known as 'MoMa'), and decides she wants to go see it. After thinking a little, she remembers where MoMa is and goes to 53rd street to see the exhibit. Here, ACDC rightly point out that there is a commonly accepted sense in which Inga believed that the museum was on 53rd street, despite the fact that she was not entertaining thoughts to that effect before hearing that there was an exhibit at MoMa:

> It seems clear that Inga believes that the museum is on 53rd Street, and that she believed this even before she consulted her memory. It was not previously an occurrent belief, but then neither are most of our beliefs. The belief was sitting somewhere in memory, waiting to be accessed. (ACDC 1998, 12)

Things aren't so simple for poor Otto, however, since he suffers from Alzheimer's disease. To compensate for the effects of his Alzheimer's, Otto carries around a notebook in which he writes down new information that he can then look up when required. ACDC claim that Otto's notebook 'plays the role' that Inga's biological memory plays in her mental life. Otto also hears about the exhibit at MoMa, but, unlike Inga, cannot remember where it is by simply looking in his head. Rather, he looks up the address in his notebook and then heads to 53rd street.

The controversy lies in how to interpret the information in Otto's notebook. According to ACDC,

> Clearly, Otto walked to 53rd Street because he wanted to go to the museum and he believed the museum was on 53rd Street. And just as Inga had her belief even before she consulted her memory, it seems reasonable to say that Otto believed the museum was on 53rd Street even before consulting his notebook. For in relevant respects the cases are entirely analogous: the notebook plays for Otto the same role that memory plays for Inga. The information in the notebook functions just like the information constituting an ordinary non-occurrent belief; it just happens that this information lies beyond the skin. (ACDC 1998, 13)

In other words, ACDC claim that since Otto's notebook was used in a way that is functionally analogous to Inga's biological memory, in that both contained relevant information and were poised to give it to the agent when needed, it is possible to conclude that Otto's notebook contained his non-occurrent (dispositional) belief concerning the whereabouts of MoMa. So his belief was outside his head.

For Otto, information pops up and retreats into his artificial memory, just like it does for normal minds: "In both cases the information is reliably there when needed, available to consciousness and available to guide action, in just the way that we expect a belief to be" (ACDC 1998, 12). The important thing to note here is the fact that it is the function played by an external object in explaining a cognitive loop that warrants attributing it membership in a cognitive process:

> *insofar as beliefs and desires are characterized by their explanatory roles*, Otto's and Inga's cases seem to be on a par: the essential causal dynamics of the two cases mirror each other precisely. (ACDC 1998, 13, emphasis mine)

5.2.2 Support for seeing the mind as extended

While the bulk of ACDC's original article motivates TXM by analyzing aspects of the thought experiment of Otto vs Inga, there is considerable empirical support for TXM as well. Perhaps the most illustrative examples of cognitive offloading to our environment comes from inattentional blindness and change blindness paradigms.

First, consider Simons & Levin' (1997) studies of change blindness. In a typical trial, a stranger engages a test subject in a conversation on the street, when suddenly two people carrying a large screen pass between the two interlocutors. Unbeknownst to the test subject, during the brief period of time in which the screen blocked the subject's visual access to the stranger, the stranger is actually replaced, mid-conversation, with another person. The fact that many subjects fail to notice the fact that they are now talking to *an entirely different person* illustrates how impoverished our representation of the visual scene actually is when compared to models that postulate rich inner-maps of our environment. Similar cases of change blindness are also strikingly displayed in studies where participants are asked to focus on a central point of a visual scene while aspects of the scene like colors of objects and presence of objects vary without the participant noticing any change. Surprisingly obvious changes can even be made within the focal point without participants noticing.[118]

Similarly striking results come from the study of inattentional blindness. In a famous study (Mack & Rock 1998), researchers asked participants to count the number of times a basketball was passed in a group of people. While the participants were busy counting passes, a person wearing a full gorilla suit walked among the passers. Surprisingly, a large proportion of participants failed to notice the person wearing the gorilla suit. This seems to illustrate ACDC's claim that "What really counts is that the information is easily available when the subject needs it" (ACDC 1998, 15).

These studies show that while we may think that vision involves creating a high definition representation of our environment, rich in detail and information, in reality most of the information is still out there – it's just that we have instructions on how to get it, if needed. This suggests that we do not store such rich maps of the outside world. Instead, we store just enough information for our current purposes and keep tabs on

---

[118] See Simons & Rensick 2005 for a discussion of these results.

how to get more information when the time comes.

Though there are real-world versions of Otto, including Patrick Jones, a chaplain whose brain injury has forced him to use his mobile phone even more that Otto relies on his notebook, cases of change blindness are less extreme and much more common, so in a way they do make a stronger case for adopting an extended approach to cognition. For in such cases

> Parts of the external world— the objects in our environment—play the same role in cognitive processing as certain neural states play: both the external objects and the neural states carry information that cognitive processing makes use of when needed for reasoning, navigation, and so on. Thus, extended states appear in a substantive role in at least one important domain of human cognition: visual perception. Taken at face value, this establishes the extended view. (Rupert 2009, 4-5)

The important point here is that whether or not an object is part of a cognitive loop isn't so much a matter of where it is as what its explanatory function is: if things in the brain play a certain explanatory role in explaining cognition, then if things outside the brain can play the same role, ACDC argue, they should also be entitled to membership in cognitive processes.

5.2.3 Extended cognition and explanation

One of the implications of adopting an extended view of cognition is to accept that certain cognitive tasks can be carried out by complex loops of continuous reciprocal causation between our heads and parts of their local (and less-local) environment. Considering the important role played by extracranial components in these 'twisted tales'[119] of complex causation, it can be tempting to reflect the contribution each component of these causal feedback loops makes in explaining the outcome of the process under consideration. However, despite the undeniable importance of external artefacts, this 'inference to egalitarianism' as Clark puts it, can obscure the fact that causal importance need not be reflected in explanatory importance. On the contrary, Clark argues that this explanatory egalitarianism is a mistake.

Given that it is impossible to describe the contribution of genes or neurons to a twisted tale of complex causality without a specification of the environment in which they are active, there is a sense in which it is redundant to require explanations of phenomena to reflect the causal importance of the environment. This is because even though parts of the environment may be more causally determinant than the gene or neuron, for example, it is the presence of these components of the twisted causal tales that makes the difference between potential outcomes to the processes in which they are involved. As Clark puts it, "explanatory priority in a given context thus turns not on what factor, if any, does the greatest amount of actual work but on what makes the difference between the cases where the outcome obtains and those where it does not" (Clark 1998, 160).[120]

This means that acknowledging the importance of the environment doesn't prevent us from usefully describing an outcome as being directed or controlled by genes or neurons, since these outcomes must take place relative to a contextual baseline.[121]

---

[119] See Clark 1998 for examples of such 'twisted tales'.

[120] Pöyhönen (2014) also focuses on the importance of differential influence in relation to explanatory relevance and emphasizes that good explanations of extended systems are supposed to draw a line between explanatorily relevant factors and background conditions.

[121] For example, in an environment-gene system, the gene is the difference maker: "a gene may be 'for x' in the simple sense that it is a feature whose presence or absence is a difference that makes a systematic (usually population-level) difference to the presence (or absence) of x." (Clark 1998, 155-6) This reasoning applies to programs and neurons as well: "The extension of the line on explanatory priority to the case of neural codes and programs is immediate. Here too we should say that a neural structure or

Thus, it is acceptable to attribute a central role to certain components of complex causal loops in explaining how behavior comes about, on the assumption that we have a reliably available ecological backdrop in which the behavior develops. Sometimes, it is OK to privilege certain components in explanations, even when they don't carry the bulk of the causal burden:

> The observation that the real workload involved in bringing about some effect may be evenly spread between allegedly "privileged" factors (genes and neural events) and other influences (environmental, chemical, bodily) cannot, I conclude, in and of itself, constitute a good reason to reject the practice of treating certain factors as special: as coding for, programming, or prescribing the outcome in question. (Clark 1998,161)

Importantly for our purposes, in explanations of human behavior this contextual baseline can include culture: "our strategies have been learned and tuned against a backdrop of culture and physical and social laws and practices" (Clark 1998, 161).[122]

Extending these claims about explanatory symmetry and difference-makers to the problem of how culture can solve the gap problem, I want to say that what we need is a way of finding a difference-maker capable of explaining the development of proto-numerical skills from systems like the ANS. Given that this development can be found in both the ontogenetic and historical levels, looking for a difference-maker to bridge the gap boils down to finding answers to the following two questions:

Q1) What makes the difference between an individual that has developed proto-numerical practices and one that has not?

Q2) What makes the difference between a culture that has developed proto-numerical practices and one that has not?

---

process x codes for a behavioral outcome y, if against a normal ecological backdrop, it makes the difference with respect to the obtaining of y" (Clark 1998, 160).

[122] This will be especially relevant in chapter 6, which deals with culture and extended numerical cognition.

Armed with these terminological tools, we are now ready to evaluate whether adapting TXM to numerical cognition can help us bridge the gap.

5.3 Numerical cognition and externalism

Given that TXM challenges the traditional boundaries of cognition, its implications are far reaching and have potential impact on any discipline in which cognition makes extensive use of external supports.[123] This, of course, includes numerical cognition. In fact, one of the first examples mentioned by ACDC is that of using pen and paper to solve an arithmetical problem, and Clark's (2008) monograph opens with an anecdote in which physicist Richard Feynmann describes the paper on which he had done his work as more than a mere record of his thought process, but as an integral part of his work.

The fact that the practice of arithmetic seems to rely on mastery of complex numeration systems like the Indo-Arabic numerals makes it look like an obvious case of cognition necessarily relying on things outside the head. This may explain why among the many explanations that have been put forward to bridge the gap between the approximate and limited output of evolutionarily-inherited systems for quantification and the precision and scope of natural numbers, most adopt an extended approach to cognition. Implicitly or explicitly, answers to the gap problem have generally relied on attributing a constitutive, irreplaceable role to extracranial objects, artefacts, or symbols in

[123] For example, if my mind includes things outside my head, is it morally objectionable to tamper with the objects to which my mind is coupled? Could deleting someone's hard drive be the equivalent of removing part of their brain? For a recent discussion of implications of TXM in ethics, see Carter & Palermos 2016.

explaining how we move beyond the limits of systems like the ANS and the OFS.[124]

We can see how externalists wish to apply TXM to numerical cognition. Here, the idea is that external representations of numbers – whether in the form of body parts, tally systems, or numerals – are constitutive parts of extended cognition systems in which numerical symbols play an essential part. As De Cruz puts it,

> External symbolic representations of natural numbers are not merely converted into an inner code; they remain an important and irreducible part of our numerical cognition…During cognitive development, the structure of the brain is adapted to the external media that represent natural numbers in the culture where one is raised. In this way, the interaction between internal cognitive resources and external media is not a one-way traffic but an intricate bidirectional process: we do not just endow external media with numerical meaning, without them we would not be able to represent cardinalities exactly. (De Cruz 2008, 487)

As we grow up, our brain adapts to the dependable presence of external numerical symbols in our environment and forms couplings with these. In a very real sense, then, our brains are wired in a way that is reflective of the numeral-enriched environment in which we grew up.

Evidence for the constitutivity of external objects and symbols can be seen in the fact that cultural variability in external artefacts affects individuals' performance in various mathematical tasks (De Cruz et al. 2010) as well as which parts of the brain are used in certain operations. For example, as mentioned in section 3.3.2, Tang et al. (2006) obtained data that seem to indicate that how we learn arithmetical practices is reflected in which parts of our brain are recruited for these practices, which suggests that the manipulation of external objects and symbols is an essential aspect of arithmetical practices. However, the fact that this study involved participants that had already

---

[124] E.g. Hurford 1987; Dehaene 2011; Lakoff and Núñez 2000; Wiese 2004; De Cruz 2008; Carey 2011; Coolidge and Overmann 2012; Menary 2015; Malafouris 2010; Ansari 2008.

bridged the gap and mastered complex numeration systems like the Indo-Arabic numerals and the abacus suggests that, at most, it shows that external artefacts are necessary to manipulate representations of exact quantities beyond an initial segment of the natural numbers, since the participants in these studies had already mastered numeration systems capable of representing quantities that dwarf the limits of unaided memory. If this is true, then it need not constitute a counterexample to the claim that the *origins* of numerical cognition do not require external support, since these origins concern an initial segment of the natural numbers that may not require cognitive offloading in order to be cognitively accessible. The idea here is that the use of external artefacts is only constitutive of numerical cognition *beyond an initial segment of the natural numbers*.

In order to assess this possibility and determine whether or not adopting a variant of active externalism can help accounts like Carey's explain the development of numerical cognition in a numeral-free world, it will be useful to characterize in which sense external artefacts can be considered *constitutive* of numerical cognition. To do this, I propose taking a closer look the notion of constitutivity as externalists apply to the practice of arithmetic.

5.4 Constitutivity of external objects in arithmetical practices

5.4.1 The synchronic level mathematical practice

Catarina Dutilh Novaes (2013) usefully frames her discussion of constitutivity by dividing it into three levels of mathematical practice: [125]

---

[125] While the fact that Dutilh Novaes' analysis of constitutivity is meant to apply to mathematics means that it goes beyond our interests, which are limited to arithmetic, most of it focuses on arithmetic, and

> the synchronic level of a person 'doing math' at a given point in time; the diachronic, *developmental* level of how an individual learns mathematics; and the diachronic, *historical* level of the development of mathematics as a discipline through time. (Dutilh Novaes 2013, 45)

Starting with the synchronic level, it does indeed seem trivially true that the practice of mathematics usually involves some kind of external object: when people perform arithmetic operations, they typically use calculators, a pen and paper, or an abacus. However, as Dutilh Novaes points out, this may simply be due to the fact that this activity is much easier when done with a little help from the world. External objects, in this sense, could merely be extremely convenient, but not constitutive of the practice of arithmetic. But constitutivity seems stronger than this: if external artefacts and/or symbol systems are constitutive of mathematical practice, then it should be *impossible* for someone to practice mathematics without these. Thus, echoing the claims made by Feynmann mentioned above, Dutilh Novaes writes that

> the claim that writing is constitutive of mathematical reasoning entails that it not only *records* independent processes; writing is in fact viewed as an integral part (embodiment) of these very cognitive processes. (Dutilh Novaes 2013, 50)

However, as Dutilh Novaes points out, this would still be a weaker sense of constitutivity, since the objects in such cases could merely be constitutive when in use, as opposed to the stronger sense of necessary for mathematical reasoning.

Dutilh Novaes claims that the degree to which we need material support for the synchronic practice of mathematics – even when it is done in our head – is an empirical matter. On this issue, Dutilh Novaes displays a strong commitment to constitutivity. The claim here is that even when we do carry out mathematical operations in our heads, these still rely on external objects, since the thoughts we have in such cogitations come

---

those parts that do not still have implications for how we see constitutivity of external support in arithmetical practices, as explained below.

from internalization of previous manipulation of external objects or symbols:

> even when a given person is apparently not manipulating symbols, such as a mental calculator, she in fact typically relies extensively on internalized versions of external devices (*at least in most cases*). (Dutilh Novaes 2013,51, emphasis mine)

To support this claim, Dutilh Novaes cites Tang et al.'s (2006) study, mentioned above, which reveals that how we learn calculate will be reflected in which part of our brain is activated when we calculate. She also remarks that the system in which a person was trained will influence the mental calculation method they use, as evidenced by the fact that notation-specific effects will be seen in their performance of various types of calculations. For example, people are typically faster in performing additions where one addend is a multiple of ten and the other is a single-digit numeral, since the task in such cases is simplified by replacing the zero with the single digit numeral. This is an effect of the place-value structure of the Indo-Arabic numeration system. Similarly, there is evidence that the practice of mental abacus, with which some people can accomplish incredible feats of calculation, is affected by motor interference, rather than verbal interference, reflecting the learning stages in which users were trained using a physical abacus (Barner et al. 2016).

From such data showing different ways in which internalization of various calculation styles affects mental reckoning, Dutilh Novaes concludes that "*at least in most cases*, mental calculations, even when performed by calculating prodigies, are essentially internal manipulations of previously mastered external symbolic systems" (Dutilh Novaes 2013:52, emphasis mine). These data show that the practice of arithmetic affects different people in different ways, based on the culture in which their mathematical skills develop. Thus, to get a better picture of the extent to which such practices rely on things outside our heads, we need to look at the diachronic level of ontogeny, given that there is evidence that the constitutivity of external support for the synchronic practice of arithmetic depends on the diachronic level of ontogenetic development of mathematical abilities.

5.4.2 The ontogenetic development of mathematical abilities

Dutilh Novaes' discussion of the diachronic, ontogenetic level of mathematical practice rightly focuses on the question of how much mathematical content is innate, and how much is learned. Dutilh Novaes points out that the ontogenetic development of numerical cognition seems to require some form of counting routine, either in linguistic or non-linguistic format, in which labels are attached to exact quantities of objects.[126] Instead of number words, many cultures used body-part-based numeration systems as their external support for numerical cognition (Dehaene 1997/2011). Also, Dutilh Novaes interprets the absence of number words for precise quantities in anumerate cultures like the Pirahã (Gordon 2004; Frank et al. 2008) and the Mundurucu (Pica et al. 2004; Pica & Lecompte 2008; Izard et al. 2008) as evidence that external symbols for numbers are needed to develop exact numerical cognition.[127] Like most authors writing on these and other cultures with limited numerical skills (e.g. Butterworth et al. 2008; Butterworth & Reeves 2008), Dutilh Novaes takes this as evidence that although language may not be necessary for the ontogenetic development of representations of natural numbers,

> exact numerical cognition is *external-symbol-dependent*; *it presupposes the very concept of exact quantities*, which may only emerge by means of explicit association to external symbols and the practice of counting beyond very small amounts (arguably, up to three). (Dutilh Novaes 2013, 54, emphasis mine)

In support of this, Dutilh Novaes cites evidence suggesting that children that have yet

---

[126] De Cruz 2008, among many others, shares this view.

[127] These cultures are discussed in more detail below.

to master counting principles and members of anumerate cultures both represent numerical distance on the mental number line (section 2.2.5) as compressed, with larger numbers increasingly densely packed towards the right. A number of experiments have indeed shown that, prior to numerical education, there is a tendency to represent numerical distance as being larger for small numerosities and smaller for larger quantities – in essence, the evidence shows a logarithmic compression of number representations before extensive numerical education (Dehaene 1997/2011). Only once an individual has mastered some formal arithmetical practices does she then represent numerical distance as being equivalent for small and large numbers. This shift in numerical distance that accompanies mastery of counting routines is interpreted as showing that, without external support and mastery of some form of counting routine, our representation of numerical distance remains tied to its innate structure. As seen in the previous passage from Dutilh Novaes, this innate structure, at best, only includes concepts for exact quantities smaller than four.[128]

We saw above that Dutilh Novaes has made a case that at the synchronic level, most cases of practicing mathematics involve internalized external symbols. But even if this weren't true of all cases, the available developmental and anthropological data concerning the ontogenetic-diachronic level is deemed more decisive:

> from a diachronic, developmental point of view, external symbols appear to be a necessary condition for the emergence of mathematical concepts and mathematical reasoning. (Dutilh Novaes 2013, 55)

---

[128] As Dutilh Novaes mentions, this evidence is disputed, since children's increasingly regular division of the mental number line could be related to their mastery of one- and two-digit numerals (Moeller et al. 2009) or to more general proportional reasoning and spatial reasoning abilities (Möhring et al. 2018). See also Núñez (2011) and Núñez et al. (2012), who have challenged the universality of this mental number line and its initially logarithmic format.

5.4.3 The historical development of mathematics

Assuming that *some* kind of external support is needed for the practice of mathematics, Dutilh Novaes argues that cases taken from the historical development of mathematical practices provide evidence for the claim that some mathematical practices are constituted by manipulations of symbols in *formal* languages. While our concern lies with arithmetic, it will be useful to consider her arguments here, as they reflect some of the commitments of externalist accounts of numerical cognition, and how they frame the origin of the content associated with external symbols. Of particular importance here is the claim that external notation is in some cases responsible for the development of novel practices.

Dutilh Novaes illustrates the role played by external notation in the historical development of mathematical practices with an example taken from Landy & Goldstone (2007), who describe Descartes' introduction of the convention that letters at the beginning of the alphabet denote constants while letters close to the end denote variables. Due to this convention, our memory no longer had to dedicate resources to which letter meant what type of object, thus allowing us to invest it in other tasks. Dutilh Novaes is right to point out that the history of mathematics is filled with such examples of major breakthroughs being 'accompanied' by improvements in notation (e.g. Indo-Arabic numerals, algebraic notations, calculus, etc.). But what form does this 'accompaniment' take? That is, do the symbols come first, or is it the mental content?

It would almost seem to be a platitude to observe that when introducing new words or symbols into the world, we introduce these to refer to something that has been discovered or observed: words label things. In the case of mathematics, it would also appear uncontroversial that introducing mathematical notation happens *after* the content associated with that notation has emerged. As Macbeth put it,

> The systems of written signs that have been devised for mathematics were devised for mathematics that already existed; it would be impossible to design a notation for mathematics without knowing at least some of the mathematics that the notation was designed to capture. (Macbeth 2013, 29, quoted in Dutilh Novaes 2013, 56)

And yet, Dutilh Novaes argues, there are cases in the history of mathematics where notation was introduced whose precise meaning was not established:

> There seem to be a number of examples in the history of mathematics where specific notations were adopted even before it became clear which concept(s), if any, they singled out. (Dutilh Novaes 2013, 56)

If this is true, then this could be evidence against a 'documentist' interpretation of mathematical practice, according to which external symbols and artefacts come *after* the internal process that has constructed the content they are associated with.

Consider the cases considered as counter-examples to 'documentism' by Dutilh Novaes, the emergence of zero and the decimal place-value system. The concept of zero itself emerged as a result of the way the place value system works, a useful place-holder to distinguish numbers like 307 from 37 and 30007, for example. The important point to consider here is that the *exact* status of zero *as a number* was only addressed long *after* its introduction in mathematical practices (Seife 2000). Thus, mathematicians were carrying out operations using the symbol for zero without knowing its exact status *as a number*. For Dutilh Novaes, this seems like it could be a case of people manipulating symbols within a system without knowing what concept that symbol refers to. After all, mathematicians used the symbol for zero for a long time before gradually coming to see it as a number:

> until the beginning of modern times in Europe, zero was not viewed as a number on a par with other numbers; instead, it was viewed as a 'gap', but this did not prevent mathematicians and users of mathematics to calculate with the symbol as if it was a number. (Dutilh Novaes 2013, 57)

The second example Dutilh Novaes wishes to levy against Macbeth's documentism is that of the introduction of the notation for calculus. The ontological status of infinitesimals was not well established for a long time, as testified by Leibniz' remark that "there is no need to let mathematical analysis depend on metaphysical controversies" (cited in Dutilh Novaes 2013, 57). According to Dutilh Novaes, the fact that people managed to use this notation without knowing exactly what it referred to means that "the concepts of infinitely small and infinitely large numbers essentially emerged from the very formalism" (Dutilh Novaes 2013, 57). Generalizing this thought, De Cruz writes that

> mathematical objects do not exist prior to the introduction of symbols that denote them, but only after notational systems have been developed that enable their expression, and that are codified by the mathematical community. (De Cruz 2007, 271)

We have just seen that Dutilh Novaes has argued that many mathematical practices – including arithmetic – are constituted by the manipulation of external symbols, and that, at least for more advanced mathematics, some of these practices require specialized notation. Dutilh Novaes wants to use the case studies of zero and the development of notation for calculus to show that there are cases where notation allows novel content to develop. This in turn supports an externalist interpretation of the role of external artefacts in the ontogenetic development of numerical abilities in both children and anumerate cultures.

Thus, if I want to show the limitations of the externalist approach to numerical cognition, it is important first that I challenge Dutilh Novaes' externalist interpretation of these historical cases, given that they are used as evidence for the claim that notation can cause the emergence of novel content in ontogeny. In the next section, I tackle this historical challenge.

5.5 Challenges to constitutivity

5.5.1 Notation and construction

The first challenge for Dutilh Novaes' views on the constitutivity of external symbols comes from her discussion of historical cases that are meant to show it is false that "the development of new notational techniques is the *consequence* rather than the *cause* of progress in mathematics." (Dutilh Novaes 2013, 56)

There are a few claims to disentangle here. First is the claim that there are cases of purely syntactic, 'de-semanticised' manipulation of symbols.[129] This claim seems uncontroversial, assuming that understanding the rules of manipulation for these symbols does not count as part of their content. Another is the claim that there are cases where we manipulate symbols without knowing exactly what they refer to. This again seems uncontroversial, given that the meaning of terms often evolves over time as we discover more things about their referent. Another claim is that in cases of ambiguous or incomplete understanding of the meaning of symbols, like zero and the notation for calculus, the notation allows us to discover things we hadn't already thought about, and thus to make progress in mathematics.

This claim is meant to count as a counter-example to Macbeth's documentism which, recall, claims that "The systems of written signs that have been devised for mathematics were devised for mathematics that already existed" (Macbeth 2013, 29). There are two interpretations of 'mathematics' here that appear relevant. On one interpretation, mathematics is a purely formal manipulation of symbols in which no meaning is required. This possibility is accepted by both Macbeth and Dutilh Novaes's

---

[129] Macbeth (2013) uses the term 'de-semantification' to describe cases where if we are equipped with the formal rules for manipulating symbols, we can learn how to follow these rules even if we do not know what the symbols refer to.

endorsement of the notion of 'de-semantification'. In this sense, however, Dutilh Novaes's claim that this could be a case of notation being adopted before it became clear what the concepts singled out were would be false, since there would be no concepts singled out, apart from the formal rules themselves, which would be explicitly laid out from the get-go. That is, if the development of notation for zero and calculus are meant to show that there are cases where "specific notations were adopted even before it became clear which concept(s), if any, they singled out" (Dutilh Novaes 2013, 56) then the possibility of de-semantification seems off the mark, given that purely syntactic manipulation of symbols need not carry with it any knowledge of what the symbols means.

On a second interpretation of this claim, the idea is that 'mathematics' refers to both the formal rules and the meaning of the symbols. It looks like Dutilh Novaes here wants to argue that symbols like '0' and those for the practice of calculus referred to things that were either not well understood or whose ontological status was ambiguous, and that these concepts were better understood as a result of the change in notation. This is considered a case of "mathematical concepts actually emerging from formalisms and notational conventions." (Dutilh Novaes 2013, 57-58) But the conclusion does not seem to follow, for a few reasons.

Consider zero first. If what Dutilh Novaes says is true, then the fact that zero came to be used as a placeholder allowed the development of content that was not available prior to this practice. Presumably, this has something to do with considering zero as a number instead of a placeholder in calculation. But what is the novel content associated with seeing zero as a number? The notion of absence of quantity, for example, was already available. As De Cruz points out, the concept of zero derived from Jain cosmological ideas of emptiness: "The emergence of zero as numerical concept in Jain mathematics was possible because it could free-ride on cosmological and philosophical concepts" (De Cruz 2007, 221). In this sense, the content associated with the symbol was already available elsewhere, and the content associated with its status as a number need not be seen as a consequence of its use in such formalism.

Further, it would seem false to claim that its inclusion in mathematical practices allowed zero to gain credibility as a number, since the peculiar status of zero as a number remains to this day, as evidenced by the fact that it is not learned in the same way as other numbers (De Cruz 2007) nor does it share many of the properties of other numbers (Seife 2000). In short, the fact that a symbol came to be used to refer to absence of quantity does not seem to justify saying that the concept of absence of quantity, or whatever zero actually refers to, was only discovered by manipulating symbols for zero.

Similarly, as Dutilh Novaes herself points out, Leibniz was aware of the ambiguous ontological status of the referents of the symbols he proposed. This suggests that the concepts that Dutilh Novaes would like us to believe came to light *via* use of this notation were already formed in Leibniz (and Newton's) head before their use started to spread, even though their precise referent had not been forged. Dutilh Novaes' claim seems to be that it is only *because* the use of such notation that their ontological status was settled. But there are a few reasons to doubt this. To begin with, it is doubtful that the ontological status of infinitesimals or zero is a settled matter. Indeed, we need not look further than natural numbers if we wish to speak of symbols whose referents are not clearly understood, given that the foundational crisis at the beginning of the 20th century revolved around questions like what sort of things numbers are.

The ontological status of infinitesimals had been problematic long before Leibniz and Newton used it in their development of differential calculus, and remained problematic after the introduction of their notation. This is testified by that fact that Eudoxus banished Democritus' atoms around 350 BC, and, much later, by the fact that infinitesimals were abandoned in the 19th century in favor of the concept of a limit, after having been described as "ghosts of departed quantities" by Berkeley and as "cholera-bacilli" infecting mathematics by Cantor (Bell 2018). This would seem to suggest both that the concepts that were behind Leibniz' introduction of his notation were already available, along with their imperfections, and that the use of this notation did not improve our understanding of their ontological status in any significant way. If this is true, it is questionable to which extent Dutilh Novaes' claim that the notation allowed its users to think of content unattainable without such notation is true, both for zero and the infinitesimals.

Perhaps we can weaken the contribution of symbols and merely say that the notation allowed us to make *progress* towards identifying what sort of things the symbols referred to, rather than settling ontological issues. Here too, in order for Dutilh Novaes' claim to be true, this would mean that the notation would allow us to think things about the content of the symbol that could not have occurred without it. But in the cases considered, what evidence do we have for this? The concept of the infinitely small was debated for millennia before the invention of the calculus and is still causing headaches for mathematicians and philosophers.

It may be worth mentioning that for this symbols-first approach to be true, we might expect the symbols themselves somehow embody the notion of infinitely small in a way that allows novel content to emerge from it. But, given the abstract nature of the mathematical infinite, it is doubtful that any particular symbol would allow such novel content to emerge. Symbols, after all, are finite.

To summarize, there is reason to doubt that the use of novel notation like the ones discussed above can be considered causes of the development of novel mental content that was not previously constructed without the help of such notation, which leaves open the possibility that the historical development of mathematics follows a documentist progression in which notation labels previously constructed content.

5.5.2 Challenges to constitutivity in ontogeny

Two other problems with Dutilh Novaes' externalism warrant discussion, both of which concern the diachronic, ontogenetic level of numerical cognition. We saw that Dutilh Novaes claims that the ontogenetic development of exact numerical content for quantities larger than three requires external support. The first problem associated with relying on external artefacts for the origins of exact numerical content beyond the subitizing range is that it does not seem to allow us to explain what happens when children learn to master quantification practices like understanding the meaning of the terms in their counting routines. In relation to the discussion of difference-makers above (section 5.2.3), if we want to explain what sets a subset-knower apart from a cardinal principle knower it doesn't look like we can look outside the head, since these children all have access to the same external support from which to learn. If the symbols or artefacts are constitutive, then it is difficult to see how we can distinguish between those individuals that have been exposed to the symbols and have mastered the novel content, and those who have not. After all, the difference between subset-knowers and cardinal-principle knowers is not outside the head, it involves some kind of induction or realization on the part of the child, as seen in chapter 4. If this insight is reached with the same access to the same external resources, it looks like the important part of this developmental puzzle is inside the head, and that external symbols are, at best, catalysts.

This seems to reflect an important aspect of learning to manipulate exact quantities: in order to say that a child (or adult) knows how to manipulate quantities, their behavior has to be guided by an *understanding* of general principles that allows them to distinguish between collections based on the number of things they contain. This understanding seems like it happens inside our head. For example, in a hypothetical case in which a person's numerical abilities were limited to asking someone else to calculate the answer to a problem and repeating it back, there would be a sense in which we would not want to attribute them mastery of numeration practices, since their behavior does not display any such understanding, even though they can 'produce' the right answer to an arithmetical problem. Rather, such a situation displays blind, de-semanticised formal manipulation of the sort discussed in Searle's Chinese room. If this is true, then externalism seems like it might have difficulty accounting for the insight that comes with realizing how counting routines work.

A separate issue concerns the potential origins of numerical content in ontogeny. I mentioned that one of the ways Dutilh Novaes uses developmental data is to justify the claim that synchronic mathematical practice involves internalization of the manipulation of external symbols learned in ontogeny. The idea is that if we have data that learning to manipulate precise quantities involves relying on external support, then any case of apparent support-free practices should be due to internalization of these practices. We may be tempted to use Carey's bootstrap as evidence of this externalist learning process, given that the first step in such a bootstrap is the memorization of meaningless number words. However, this would be leaving aside one of the most distinctive elements of Carey's approach, which is that these symbols are mapped to the representations of set-based quantification. The content of set-based quantification is evolutionarily-inherited and serves to ground the initially meaningless list. This means that such content is available to us without any external support. If this is true, then, as Dutilh Novaes accepts, there can be some primitive and restricted range in which we could manipulate exact quantities – perhaps, for example, ONE versus MANY.

Here, as in other levels of analysis of the origins of mathematical practices, we are confronted with a 'chicken-egg' problem, as Dutilh Novaes notes. Take the case of anumerate cultures like the Pirahã and the Mundurucu: is the fact that their languages do not contain number words the reason for their lack of numerical practices, or is this lack of practices the reason that their language lacks number words? The answer to this question has important consequences for externalism about numerical cognition, given that if it were to turn out that the latter option are true, then we could argue that we do not need external linguistic support to develop representations for exact quantities. The case of anumerate cultures is not settled yet, though Dutilh Novaes' remarks that such cultures could be considered as having labels for exact quantities up to three implies that *some* restricted numerical content would be available without external support. If this is the case, then constitutivity needs to be qualified, since the fact that *some* numerical content doesn't require external support for its development begs the question of *how much* such content can be reached without attributing a constitutive role to external support. To explore this issue, I take a look at the origins of non-linguistic counting routines in the next section.

5.5.3 Which came first, the number or the numeral?

As mentioned above, there is consensus that the ontogeny of representations of exact numerical quantities requires a counting routine to go through. When asking where such counting routines come from, Dutilh Novaes seems to think that their origin depends on having representations of discrete quantities:

> humans have a long history of developing calculating devices/objects such as counting rods and abacuses, and each of them presupposes that quantities be represented so that they can be 'operated on' for calculation. (Dutilh Novaes 2013, 51)

In order for this to support externalism, it must be possible to come up with an externalist explanation for the origins of representations of the notion of precise quantity, insofar as the development of counting devices is said to depend on such representations. However, the origins of a representation of the notion of precise quantity itself seems to rely on an important insight of a similar nature to the one that occurs when we bridge the gap, since none of the core cognition modules has explicit representations of discrete quantities, let alone of the notion of precise quantity itself. The fact that Dutilh Novaes seems to agree that it is possible to form such a representation *before* we have counting routines suggests that the important notion of precise quantity can develop without the benefit of the external support used in counting routines. If this were true, then there is reason to doubt that we need external support in the strong, constitutive sense of the term, given that the representation of discrete quantity is achievable without such support. While there is no doubt that external support may facilitate the acquisition of this content and extend its domain of application, these would appear to be weaker senses of constitutivity that make room for the possibility of internalist development of exact numerical quantity.

Perhaps I am misinterpreting Dutilh Novaes' claim here, but there are good reasons to think this misinterpretation happens to be true. Consider the opposite situation, where counting requires external artefacts for its acquisition. As Dutilh Novaes observes, at least in ontogeny, "Counting is a practice which only emerges upon explicit instruction" (Dutilh Novaes 2013, 54). On pain of a regress of explicit instructors, however, it must be possible to come up with counting routines on one's own, at least, in a primitive form, such as having labelled a few precise quantities and understanding that they are labels *of* precise quantities. If this is true, then there is reason to think that external supports are not *required* to develop such practices. They are merely extremely convenient. And if this is true, then the associated claim that any synchronous case of internal computation is an internalized version of external practices is weakened.

Even if these three problems turned out not to be a worry for the externalist, Dutilh Novaes' argument for the constitutivity of external symbols and objects in the synchronic practice of arithmetic is based in large part on models of the ontogeny of exact numerical abilities like Dehaene's and Carey's that crucially rely on external artefacts. However, as I mentioned earlier in the description of my origins problem, the origins of these externalized practices may be a problem for the externalist, since these models' reliance on external symbols for numbers in ontogeny begs the question of where these symbols come from *qua symbols for precise quantities*, given that using objects as symbols for precise quantities must have been the creation of human individuals. Thus, it seems problematic at this point to consider these as evidence that external symbols are necessary in ontogeny, given that we still haven't figured out how the first external representations of precise quantities emerged in a world where numerals and number words were absent. While the precise details of such a story will probably always remain a matter of speculation, it would still be useful to get a better idea of how this could have happened. To this end, in the next section I discuss data on the development of numeration practices in the Near-East.

5.6. A case study of the historical development of numerical cognition

5.6.1 Material Engagement Theory

For Lambros Malafouris, to explain "the ontogenetic and phylogenetic passage from approximate to exact arithmetic" (Malafouris 2010, 35), it is necessary to look beyond neurological data concerning current-day arithmetical practices. This is because explaining this passage involves explaining how "the sapient mind make that leap forward, overcoming the limits of approximate numerical thinking" (Malafouris 2010, 35), which, at least initially, occurs when "an explicit vocabulary of number words does not exist" (Malafouris 2010, 35). Thus, Malafouris expresses doubts concerning the ability of language-based accounts of delivering an answer to how we first bridged the gap:

> I do not dispute the fact that it is language competence that enables above all the development of the verbal symbolic number system essential for the development of exact calculation and higher mathematics. What I argue is that

> although this may well be the case when verbal number exists in the Vygotskian developmental 'zone of proximal development', it cannot be used to account for the development of numerical thinking in a context where such verbal numerical competence does not yet exist. (Malafouris 2010, 38)

However, while this may make Malafouris look like he shares the worries expressed by my origins problems, and thus that he would be inclined to reject adopting an externalist account of numerical cognition, this is not the case.

This is because Malafouris' also adopts a brand of externalism, which he dubs *Material Engagement Theory* (2010, 2013). Material engagement theory can escape some of the problems associated with explaining the origins of extracranial objects with numerical meaning, because it focuses on the role of the body in shaping the development of our numerical abilities.[130] Instead of focusing on neurological data concerning the way arithmetic is currently practiced, Malafouris focuses on data from cognitive neuroscience that deals with non-linguistic numerical processing, as well as with data pooled from archeological and anthropological research, to try to piece together the conditions in which the initial origins of numerical practices occurred.

For example, as evidence of the influence of culture on numerical cognition, Malafouris notes the evidence mentioned above (section 2.2.5) that the Space Number Association Response Code (SNARC) can be influenced by cultural practices, including the direction in which we learn to read (Zebian 2005). This shows that culture has a profound effect on how we process numerical information, including the relationship between number and space. Malafouris also mentions the importance of counting routines for the development of arithmetically-viable numerical content, as mentioned by Pica et al. (2004) as a possible reason for why the some anumerate cultures lack

---

[130] In this sense, Malafouris' brand of externalism can be seen as adopting aspects of both enactivism (Varela et al. 1991) and distributed cognition (Hutchins 1995).

representations of exact quantities beyond the subitizing range.

Malafouris, like Dutilh Novaes, argues that counting routines are not necessarily verbal, and can include tallying systems and body-part systems: "a system of counting words may not necessarily be the only system of counting" (Malafouris 2010, 38). For Malafouris, such cases of non-verbal counting routines are examples of "an embodied sensorimotor association of quantity with the movements of touching various parts of the body in a fixed order" (Malafouris 2010, 38). Such cases involve the substitution of abstract numbers for "a complex enactive bodily system" (Malafouris 2010, 38), where we can express specific dates, for example, in terms of numbers of days following a full moon, by pointing to parts of our bodies. Thus, in looking for the non-verbal origins of representations of exact quantities, Malafouris highlights the fact that many objects in our environment can be used to count things, including knotted strings, sticks, pebbles, and bones.

As evidence for the variety of tools that can be used to count things, the archeological literature has turned up an impressive number of relics of ancient counting systems, including bones with carved notches in them dating back around 30 000 years. For example, the notched bones from Abri Cellier in France date back at least 28 000 years (Overmann 2014), while Ifrah (1998,123) claims notched sticks were used as tallies at least 40 000 years ago. Ifrah also notes that one wolf bone dating back to around 30 000 years contained 55 markings grouped in series of five, which, he claims, "demonstrates that at that time human beings were not only able to conceive of number in the abstract sense, but also to represent number with respect to a base." (Ifrah 1998, 119-120). Knotted string was used in many independent civilizations for accounting purposes (Ifrah 1998), while hands also appear to have been used for accounting and calculating purposes, as evidenced by 27 000-year-old hand stencils found in caves in France (Overmann 2014).

5.6.2 Tokens for numbers

According to Malafouris, the first evidence we have of representations sharing the abstract properties of natural numbers is that found in clay tablets dating back around 5000 years to the Sumerian empire that appear to have been used to keep records of commercial transactions (Ifrah 1998; Schmandt-Besserat 2010). The inscriptions on these tablets were based on earlier clay tokens used by the same culture around 7500 years ago. Initially, these tokens took one of six shapes (cones, spheres, cylinders, discs, tetrahedrons, and ovoids), each of which represented a particular type of object. According to Schmandt-Besserat, these tokens show that counting and accounting started when agriculture and commercial trade required better quantification methods. Tokens then took on a much larger variety of sizes to represent the increase in variety of traded goods. In such accounting systems, quantities of objects were represented by the tokens using one-to-one correspondence (hereafter: 11C) between tokens and quantities of traded commodities.

When the need to keep track of exchanges became more complex, the Sumerians invented clay ball-shaped 'envelopes' in which the tokens would be kept as records. To mark the content of the envelopes, the tokens were imprinted on its sides, so that their contents could be inspected without needing to break the envelope (See Figure 5.1). This marks a transition from three-dimensional token to two-dimensional tokens. These envelopes were eventually replaced with tablets on which the shape of the token was impressed, still relying on 11C between numbers of goods and (impressed) symbols for these. Then, in a later stage, the tokens were drawn or carved in the clay, instead of being the result of imprinting.

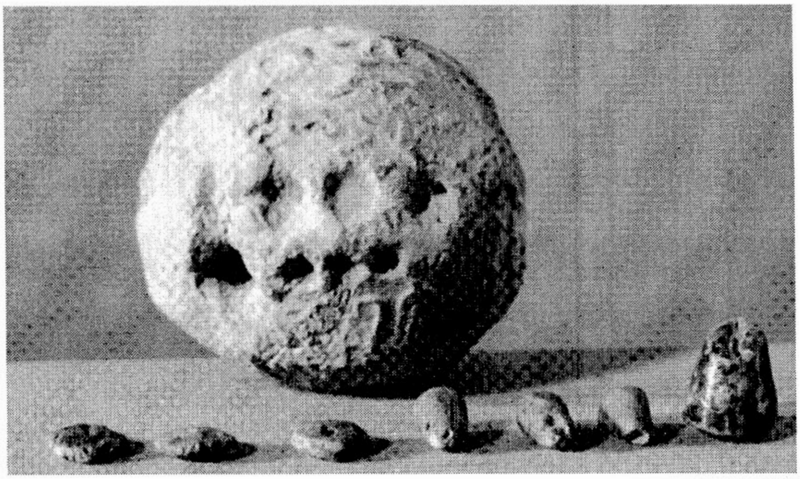

Figure 5.1. Clay envelope with impressed token-shapes. (Schmandt-Besserat 1996)

The important moment in the development of abstract notation for number came when the accounting practice shifted away from 11C: at some point around 5000 years ago, the pictographs representing the tokens were no longer drawn the same number of times as the objects represented were counted. Rather, a single pictogram representing the object type was drawn, along with notches of different sizes next to it to represent how many of the objects were represented. For example, a large wedge represented sixty units, a small wedge represented one unit, and a circular imprint represented 10 units (see Figure 5.2 for an example of such abstract notation). As Schmandt-Besserat notes, the increase in efficiency of such notation is remarkable: for 33 units, 7 inscriptions were now needed, instead of the 33 that would be required by 11C.

Figure 5.2. Clay tablet with abstract symbols for quantities. This tablet displays 33 units of oil, as indicated by the 3 round impressions (10 units) and three wedges (units) next to the symbol for oil, to their right. (From Schmandt-Besserat 2010)

According to Schmandt-Besserat,

> Pictography thus marks the extraordinary event when the concept of number was abstracted from that of the item counted…It was also momentous in overcoming one-to-one correspondence, which had governed counting during the entire token era (Schmandt-Besserat 2010, 31)

It should be noted that while this notation was indeed abstract, in the sense of having a symbol for a quantity other than one, symbols were still commodity-relative. For example, symbols that could stand for 10 units of animals could refer to 6 units of grain.[131]

---

[131] Beller & Bender (2008) claim that number words are one of the necessary ingredients to the development of 'complex numerical cognition', but they also claim that, contrary to popular belief, object-based terms can be more useful in calculation, given the size of their counting units, which extends the limits of counting. In this sense, the fact that Near-Eastern symbols were still commodity-relative need not count against their level of complexity or advancement, since there is evidence of some

5.6.3 Problems with Material Engagement theory

According to Malafouris, the historical development of Sumerian accounting techniques that led to the abstract representation of precise quantities of objects shows that material engagement with cultural artefacts is a necessary precursor to such abstract representation of number. Malafouris rightly points out that to answer the gap problem we have to identify "the causal role of this form of meaningful material engagement" (Malafouris 2010, 39). However, his answer to this question seems to run into a few difficulties. For example, in his explanation of how manipulating clay tokens allows the development of an abstract conception of numbers, Malafouris claims that

> the use of clay tokens provides the necessary external scaffolding, with a dynamic and constitutive role, for the emergence of arithmetic competence… counting with clay tokens should be seen as an integrative projection between mental – the basic biological approximate 'number sense' (Dehaene 1997) – and physical – for instance, fingers or clay tokens – domains of experience. It is the resulting structural coupling…that brings about the possibility of the meaningful cognitive operation we know as counting and not some innate biological capacity of the human brain. The clay tokens do not stand for numbers, as it may seem; the clay tokens bring forth the numbers and make visible and tangible the manipulation of their properties. (Malafouris 2010, 40)

One problem here is that, contrary to what would appear to be described by the notion of 'integrative projection', the ANS is not equipped to handle either 11C between tokens and physical objects, as the evidence presented in chapter 2 concerning the limited precision of this system shows. Not only is it doubtful that the ANS can underlie representations of quantities within the confines of the restricted subitizing range, but it is also doubtful that a mapping between the ANS and external symbols or artefacts

---

numeration systems going from more abstract commodity-independent format to object-specific format.

underlies the generalization that lead to 'arithmetic competence', as seen in Carey's criticism of Dehaene's mapping in chapter 3.

Moreover, Malfouris' account attributes considerable importance to the fact that we are manually manipulating numbers instead of having to think about them and claims that grabbing and manipulating tokens reinforces and builds neural connections between the ANS and areas of the brain related to manual tasks like grabbing and pointing. But while there is indeed strong evidence that hand gestures influence numerical processing,[132] the evidence against attributing a necessary role to finger counting in the development of our understanding of natural numbers is also quite decisive (Crollen et al. 2011), which suggests a similar limitation to manipulating clay tokens with our hands.

Also, while there is a sense in which grabbing a clay token is a more 'concrete' representation of numbers, there is another sense in which the fact that numbers are generally considered abstract entities means that there is nothing in the grabbing of a clay token that somehow makes the representation less abstract, given that there is nothing in the material shape of the token that somehow better embodies what numbers are, due to their abstract nature.

Perhaps more importantly, however, despite its focus on non-linguistic development, Malafouris' reliance on external tokens seems to share some of the same issues at the heart of my origins problem. For while we can agree that the tokens themselves can be the result of hands molding clay, the question of how these come to be used as symbols for precise quantities is not settled in this account. Malafouris claims that the manipulation of clay tokens makes the problem of counting easier to solve by making

---

[132] The relationship between representations of fingers and representations of numbers is receiving increasing attention in the numerical cognition literature. See Crollen et al. 2011; Fayol & Seron 2005; Moeller et al. 2012; Andres et al. 2012.

it a spatiovisual problem instead of a conceptual problem:

> the vague structure of a very difficult and inherently meaningless conceptual problem – that is, counting – by being integrated via projection with the stable material structure of the clay tokens is transformed into an easier spatiovisual problem. However, spatiovisual problems can be directly manipulated and manually resolved in real time and space. As Schmandt-Besserat observes…tokens made it possible to visualize and manipulate numerocity [*sic*]. Thus the problem – that is, counting – becomes meaningful. (Malafouris 2010, 40)

While it seems likely that visuospatial problems are indeed easier to solve in many cases, as illustrated in the case of *Tetris* mentioned above, the problem here is that the mere fact that counting is considered a problem at all presupposes cognitive resources whose origins is left unexplained on this account. As mentioned above, in order to count, we have to have a representation of precise quantity. This means we had to have the concept of precise quantity in order to use tokens as a solution to a counting problem. Representations of precise quantities arguably underlie our basic numerical abilities, given that in order to consider any two collections as distinct based on the number of items they contain, it is necessary to have the concept PRECISE QUANTITY. However, as mentioned above, neither the ANS nor any other core cognition module allows such distinctions based on explicit representations of precise quantities. While the ability to put items into 11C could arguably be innate, doing this as part of a task of precise quantification is not. To see why, in the next section, I take a closer look at the quantificational abilities of anumerate cultures like the Pirahã and the Mundurucu.

5.7 Quantificational abilities in anumerate cultures

5.7.1 The Mundurucu

To clarify the relation between language and arithmetic, Amazonian cultures like the Mundurucu and the Pirahã , whose languages lack a productive number lexicon, were extensively studied in the past fifteen years or so. Consider the Munduruçu first. The Mundurucu lack words for numbers larger than five, and even in this restricted range, their use of numerical language is inconsistent, as evidenced by the fact that they can alternately describe a collection of three objects with the Mundurucu word 'ebapug', which is most frequently used to describe collections containing three objects, but also can be used for collections of two and four objects (Izard et al. 2008). In other words, it looks like they treat even small numbers as approximate quantities. While most Mundurucu can now recite a count list of Portuguese counting words, they still use these foreign words to refer to approximate quantities, which suggests that, like the subset knowers in Wynn's studies, they have yet to come to an induction that allows them to associate the meaning of these words with precise quantities.

While they do not have a linguistic counting routine, the absence of number words does not prevent some Mundurucu from being able to use their hands and toes to match the number of dots they see, when prompted to. To determine to which extent their lack of a productive system for number words influences their performance in rudimentary numerical abilities, Pica and colleagues (2004) asked Mundurucu participants to point to the dot array that had the larger number of dots, using dot arrays containing between 20 and 80 dots. Performance was well above chance, though slightly below that of educated westerners. Distance effects were observed, suggesting that their numerical abilities are recruiting the same system as educated adults (i.e. the ANS).

These researchers also tested Mundurucu ability to perform rudimentary arithmetical operations by exposing them to images of distinct collections of dots being placed into a container and then asked them to compare its contents to a third collection of dots. Again, performance was well above chance and displayed distance effects. In a third task, researchers tested whether the Mundurucu could perform operations on exact quantities of objects. Participants were asked to use a word to describe how many objects were left after a collection of dots was subtracted from another. While the number of dots in both collections was out of range for Mundurucu number vocabulary, the result of the subtraction was in their number lexicon. Here, their performance was much lower than that of educated westerners. The authors speculate that the absence of a counting routine may explain the limitations of the Mundurucu numerical abilities, since "counting may promote a conceptual integration of approximate number representations, discrete object representations, and the verbal code" (Pica et al 2004, 503).

The most important result for our purposes is that the Mundurucu and other anumerate cultures are able to put objects into 11C. Evidence for this comes from many sources. For example, when asked to match the number of dots presented on a screen using seeds, the Mundurucu perform above chance. Similarly, consider the following task used by Pica and collaborators:

> participants were presented with an animation which involved sets of red and black puzzle pieces and a can. At the beginning of each trial, red and black pieces were presented in one-to-one correspondence, and the can was shown to be empty. The black pieces stayed in place during the whole trial, while the red pieces started to move and disappeared inside the can. At this point, a transformation occurred which affected the hidden set of red pieces: some pieces were added or subtracted, or pieces were replaced by other ones. At the end, some of the red pieces came back in front of the black pieces, and participants were invited to judge whether the box was empty or not by clicking on one of two alternatives. (Izard et al. 2008, 501)

Performance on this task was also above chance levels. This shows that the absence of a counting routine does not prevent the ability to match collections in terms of the quantity of objects they contain using 11C. However, given the variable way in which

Mundurucu words are used to denote precise quantities, it looks like their ability to put collections into 11C has not allowed them to form explicit representations of PRECISE QUANTITY. This would appear to be strong evidence that it is possible to use 11C without having a general representations of PRECISE QUANTITY, thus supporting my claim that Malafouris and Dutilh Novaes cannot rely on 11C to explain where we get representations of precise quantities from.

I take this as evidence that 11C is *not* enough to get the concept of precise quantity, since for this concept to be present, there needs to be an insight or understanding that what sets the collections apart and what is being matched is explicitly represented in terms of the exact number of items they contain, whereas all we have is evidence of here is an understanding of SAME or DIFFERENT, not an explicit representation of what could explain the difference between these collections.[133] If the notion of PRECISE QUANTITY were available to the Munducuru, even within a restricted range, then we would expect them to use the same labels to refer to collections that contain the same number of items, just like other words of their language refer to the same property of the objects being named. The fact that this is not the case shows that their ability to put objects into 11C is not guided by the task of putting objects in *11C in terms of their precise quantity*. Rather, they could simply be guided by the task of matching members of collections one by one, until this task is no longer possible, without having come to a realization that this process allows matching of collections in terms of quantities. This sets their 11C skills apart from those of Sumerian accountants, whose 11C was, according to Malafouris, aimed at solving a quantity-related problem.

This interpretation seems confirmed by one of the most prominent researchers of

[133] As seen in section 1.4, similar considerations prevent us from attributing numerical content to the OFS.

Munduruсu numerical abilities, Pierre Pica:

> We would like to suggest that the absence of exact numbers in Munduruсu is closely related to the fact that, although Munduruсus recognize sets and individuals, *they do not make sets out of individuals through distributive quantification* (Pica & Lecompte 2008, 511, emphasis mine)

This seems consistent with interpreting an absence of understanding of the concept of precise quantity, since if the Munduruсu did understand the general concept of precise quantities, then we might expect them to be able to individuate collections based on this understanding of precise quantity. And yet, the evidence seems clear that they can perform 11C, further suggesting that 11C does not entail an understanding of what precise quantities are.

In the Munduruсu language, the word often associated with collections of five objects means ‘a handful’ (Pica & Lecompte 2008). My claim is that while such words may appear to describe the possibility of putting collections of objects into 11C with our hand, this does not mean that the Munduruсu have an understanding of what the referent of this word and the referent of other terms in their restricted numerical lexicon have in common, namely, that they refer to the precise quantity of objects in collections, as individuated in terms of what these collections can be put into 11C with (e.g. parts of their hands). The situation seems analogous to that described by Wynn’s subset knowers, where children can use ‘three’ to describe collections in terms of the possibility of putting their elements in 11C with elements of words in their count list, and yet do not have an understanding that these instances of 11C have in common the fact that they are instances of the application of the concept PRECISE QUANTITY. In both cases, what is missing is an insight, a realization, that allows using labels for other instances of 11C:

> Around the age of 3, Western children exhibit an abrupt change in number processing as they suddenly realize that each count word refers to a precise quantity. This “crystallization” of discrete numbers out of an initially

> approximate continuum of numerical magnitudes does not seem to occur in the Mundurucu. (Pica et al 2004, 503)[134]

5.7.2 The Pirahã

Another extensively studied culture with limited numeration skills is the Pirahã, whose number lexicon is arguably worse than the Mundurucu. While Gordon (2004) interpreted Pirahã language as containing words for 'roughly one', 'roughly two', and 'many', on Everett's (2005) interpretation, the Pirahã words for 'one', 'two' and 'many', are better translated as 'small size/amount', 'a bit bigger size/amount', and 'cause to come together/many'. To determine which of these interpretations is better, Frank et al. (2008) tested the words used by the Pirahã to describe collections of spools of thread containing from one to ten spools. Importantly, the number of spools was presented in ether ascending or descending order (i.e. in some trials, the first stimulus contained one spool, and then each successive trial contained one more spool, while in other trials, the reverse situation was presented, starting with ten spools and descending back to one).

Unexpectedly, the Pirahã used different words in ascending and descending trials. For example, the word that could be the Pirahã equivalent of 'one' was used to describe the one-spool stimulus on ascending trials. However, on descending trials, this word

---

[134] Further support for this interpretation comes from the fact that animals are able to put objects into 11C, learn to label collections in terms of the quantities of objects they contain, and even associate symbols for these quantities. And yet, they fail to generalize their ability to novel cases. For example, Japanese primatologist Tetsuro Matsuzawa (1985, 2009) managed to teach a chimpanzee named Ai to associate various symbols with collections of objects. This chimp was taught to associate sets of objects with names, colours, and number symbols. Ai eventually managed to successfully associate the first nine digits with collections displaying the associated numerosity, and has also learned to order the numerals according to their magnitude. Since then, a few other chimps have also learned to use numerical symbols correctly. See also Boysen et al. 1996.

was used to describe collections of up to six spools. This seems to show considerable variability in the Pirahã 's labelling of collections based on their size, prompting researchers to conclude that this culture may even lack a word for EXACTLY ONE. Given that their language does not display such variability in other domains – the word for water is not used to describe fire, for example – this seems to suggest that they do not have the concept of precise quantity – or, if they do, they have no linguistic label for it. And yet, as evidenced by their performance in matching tasks, they can perform 11C on collections containing numbers of objects that go beyond the subitizing range. This again, suggests that 11C is not enough to get the concept of precise quantity, as suggested in the following interpretation of Pirahã 11C abilities:

> the one-to-one matching task itself can be completed via a simple algorithm, "put one balloon down next to one spool." At no point during the task must participants represent the cardinality of the entire set. They need only to understand that, in the application of this algorithm it is exactly one balloon that must be matched to exactly one spool. (Frank et al. 2008, 823)

5.8 Limiting the scope of constitutivity

The previous section showed that there is evidence supporting the claim that the ability to put objects into 11C does not guarantee possession of the concept PRECISE QUANTITY. This, in turn, suggests that the development of counting routines, which are considered essential to the emergence of advanced numerical cognition, cannot be explained solely by an ability to put objects into 11C. And yet, as we saw, Malafouris seems to frame the solution to the gap problem in terms of solving a counting problem by material interaction with clay tokens. But, given that none of our core conceptual modules can represent a counting problem nor the concept of precise quantity on which such a problem depends, this account, like other externalist accounts of numerical cognition, seems to take for granted resources whose origins remain unidentified.

The problem of the origins of numerical cognition is intimately related to that of the origins of representations of precise quantities. Malafouris' account seems to skip this crucial explanatory step of how we come to represent a problem whose solution involves precise quantification using 11C. For example, how could a person come to think of representing ten jars of oil with a single symbol unless that person had come to a kind of insight that allowed them to think of precise quantities in the first place? For this to happen, it would seem that someone has some kind of insight, of the same sort that allows the transition from subset-knowing to CP-knowledge. Sumerians were using 11C for their accounting purposes for a long time before symbols for quantities were eventually used. The question here is: how did this transition occur? Given that the practice of 11C was used for so long without there being symbols for precise quantities, this would seem to suggest that the ability to do 11C isn't sufficient to explain the development of abstract representations of precise quantities. Rather, what is needed is an explanation of the internal insight that led to the emergence of the concept PRECISE QUANTITY.

Similarly, given that both members of anumerate cultures and brilliant mathematicians both use their hands all the time, including to manipulate objects, it would seem necessary to identify what is different in the manipulation involved in Sumerian accounting that allows their hands, but not those of members of anumerate cultures, to lead to the development of abstract representations of precise quantities. But the difference between manipulating tokens for accounting and to throw them at an enemy, say, presumably has something to do with the task guiding the manipulation. Here, the task is one of precise quantification. I thus claim that in order to manipulate tokens with the task of precise quantification, we have to have already made significant progress in bridging the gap.

Such considerations, if they ring true, suggest that there is an essential ingredient missing from externalist accounts of numerical cognition, even when they do not rely on complex numeration systems like Indo-Arabic numerals or number words. This missing ingredient is what sort of insight allows both the systematic representation of discrete quantities and the possibility of using external symbols for these. Given that any external object that we wish to claim was used to develop representations for precise quantities – including fingers – seems to depend on the occurrence of an important realization or insight of the sort described by Carey in order to figure in a *numerical* or *quantificational* task, including counting objects or answering a parent's *how many* question, it is likely that the answer to our origins problem lies inside our heads.

Of course, there is a sense in which this origins problem need not affect the constitutivity of external support for numerical cognition in general, given that there is clear evidence that external symbols and artefacts are necessary to represent numbers beyond a rather restricted range. This seems compatible with adopting a weakened form of constitutivity, in which external artefacts and symbols are necessary for numerical cognition in domains that require cognitive offloading and memory aids.[135] While there is no reason to think that this range can be given a precise limit, the fact that external artefacts are often described as being useful memory aids and tools for cognitive offloading suggests that the upper limit on a potential external-support free

---

[135] There is evidence that this is the major role played by extracranial parts of extended numerical cognition loops. After all, the information in Otto's notebook was in his head before it got on paper, assuming he cannot write unconsciously. Similarly, De Cruz writes "External symbols can provide *anchors for thoughts* that are difficult to understand or represent" (De Cruz 2007, 249) and that writing down intermediate results in calculations "transforms what would have been a difficult cognitive task into an easier perceptual task" (De Cruz 2007, 242). A weak sense of constitutivity, compatible with attributing a facilitation role to external artefacts, also seems implied in the following passage:

> We tend to view the role of non-biological extensions to the brain as that of providing more (and more durable, shareable and stable) memory. However, external media may also be regarded as an extension of cognitive processes. Some actions performed by use of external media solve problems more easily and reliably than if they had been solved in the mind alone. The objects we use when making calculations only acquire their meaning once we have come to understand the calculation process and the meaning of the operations that define it. (De Cruz 2007, 241)

range of representation of precise quantities would be a function of a person's working memory.

Such a restricted numerical range for the constitutivity of external objects need not be considered a major problem for the externalist, however, given that there is already evidence of discontinuities in our numerical abilities depending on the external support we are using. While we can doubt that cultures with limited explicit symbols for numbers like those using body-part numeration systems can distinguish between quantities beyond their numeration system's domain of application or reason about the infinity of natural numbers, their ability to represent and label precise quantities as such still sets them apart from anumerate cultures.

As evidence of constitutivity for larger numerical domains, Schlimm (2018) argues that which levels of numerical abilities can be reached is related to the expressive power of the numeration systems with which we are familiar: "the use of numerals and other formal systems of representations is constitutive for the development of an advanced conception of numbers" (Schlimm 2018, 213). Among other skills, such an advanced conception of numbers involves more than an ability to count the items in a small collection, which would appear to be all that we could achieve with any reasonable degree of success without the use of memory aids and cognitive offloading. For example, the ability to compute arithmetical operations on numbers larger than 100 and the understanding that there is no greatest natural number both seem to require understanding the tools offered by advanced numeration systems like the Indo-Arabic numerals. Even number words are limited when compared to such systems, given every language has a limit to the quantity it can express with a single lexical item, while the recursive construction of numerals with Indo-Arabic symbols has no such limitation.

This constitutivity for larger numerical domains seems to confirm evidence mentioned above in relation to the variability of cultural invasion of brain tissue (Tang et al. 2006). The effects of the structure of a numeration system are also illustrated by the unit-decade compatibility effects, mentioned in section 2.2.2 which shows that the place-value structure of the Indo-Arabic system affects processing of multi-digit numerals, which would not happen the same way in systems with different bases, since unique symbols could represent larger or smaller quantities, depending on the base. This shows that what numeration system we use matters in terms of what we can represent and how easy it is to represent it. Such differences in the degree of expressive power of numeration systems on which we rely to complete numerical tasks seems consistent with my claim that an initial segment of the natural numbers that goes beyond the subitizing range can be attained without external support, and that to go beyond this restricted numerical range, external supports are necessary to help processing.

In this chapter, I have argued that the externalist approach to numerical cognition cannot answer my origins problem because it fails to account for what makes the difference between a person that has bridged the gap and one who hasn't, when both individuals have access to the same external resources. To account for this difference, I claim, we must look inside the head. If what I've argued above concerning the importance of internal realization for the emergence of PRECISE QUANTITY makes sense, embracing an extended cognition framework does not help solve the origins problem, since accepting external objects as constitutive parts of cognitive systems does not tell us where the content associated with these objects comes from. Coming back to Otto, while we may accept that his notebook is indeed a constitutive part of his cognitive routine, it is important to keep in mind that everything he wrote in his notebook got there because it was in his head beforehand. Applying this analogy to the case of (extended) numerical cognition seems to suggest that we must look in our heads to find out where numerical symbols and artefacts come from.

However, perhaps the externalist can reply to my origins problem by claiming that we can separate artefact-free and artefact-based cases of ontogeny, and that we can explain the difference between these via mechanisms of cultural evolution. This way, the question of the historical origin would be completely separable from the question of the ontological development of number concepts, and the constitutivity of external support would be secure. On this view, how number concepts arose in individuals in the past is different from present day ontogenesis. To determine the worth of this option, I explore the role of culture in the construction of novel numerical content in the next chapter.

CHAPTER VI

CULTURE AND NUMERICAL COGNITION

6.1 Introduction

In chapter 5, I discussed whether adopting an externalist approach to numerical cognition, in which things outside our head play a constitutive role in arithmetical practices, can help us bridge the gap between our evolutionarily-inherited cognitive systems and the systems responsible for more advanced numerical abilities. I argued that this approach fails to deliver an account of how we bridge the gap in environments where no numerical symbols or artefacts are present, and that this means the constitutivity of extracranial objects in numerical cognition does not apply to a restricted initial segment of the natural numbers. However, the externalist can try to appeal to processes of cultural evolution to explain the emergence of novel numerical content and its association with external artefacts by a gradual accumulation of cultural variation.

After all, talk of cumulative culture taps into an important intuition concerning how number concepts emerged: it is patently false to claim that these appeared fully formed, complete with all their formal properties, in a single individual's head. Rather, it seems more appropriate to describe the history of mathematics as one of individuals reflecting upon historically constructed ideas and adding their bit to an increasingly large body of knowledge. At no point does the emergence of novel mathematical content require a person to reinvent the whole body of mathematics, and the same would seem to apply to the simpler concepts involved in numerical cognition.

Given that mechanisms of cultural evolution can explain how content changes and evolves over generations, the externalist could appeal to these to explain how number concepts evolved, and we could thereby accept two (or more) ontogenetic stories to accommodate various levels of mathematical knowledge, thus disarming the bulk of my origins problem. Questions about the increasing role of numerical symbols in numerical cognition could be answered by mechanisms of cultural evolution, transmission, and inheritance (e.g., Dawkins, 1976/2006; Aunger, 2001; Richerson & Boyd, 2005; Sperber, 1996). Such mechanisms could perhaps explain how symbols in the environment could have acquired novel numerical content. For example, symbols for approximate quantities could have gradually spread and evolved due to their increasing usefulness in advancing societies.

To see whether cultural evolution can help the externalist explain how we bridge the gap by providing a means of developing external artefacts with numerical content from contexts where no such artefacts are found, this chapter explores mechanisms of cultural evolution and their relationship with extended cognition. In this chapter, I want to explore two potential replies from the externalist to my origins problem, both of which involve recourse to the role of culture in the development of numerical cognition. The first option is to explain the development of objects and symbols with numerical content by mechanisms of cultural evolution. To determine to which extent cultural evolution can help the externalist bridge the gap, I will follow the main lines of Helen De Cruz's (2007) Darwinian approach to mathematics in the first few sections of this chapter. The second option is to extend the mind further out of the head and include culture itself as part of extended cognition loops. This way, the burden of explanation does not fall squarely on an individual's head, but on groups of individuals and population-level processes. To explore this second option, I will focus on a recent externalist proposal that is particularly explicit about what culture brings to the gap problem: Richard Menary's (2015a) enculturated approach to mathematical cognition.

6.2 Helen De Cruz's Darwinian approach to numerical cognition

6.2.1 Introduction: Darwinism outside of biology

In *The Descent of Man,* Charles Darwin noted an analogy between evolution of species and the evolution of language: "The formation of different languages and of distinct species, and the proofs that both have been developed through a gradual process, are curiously parallel" (Darwin 1871/1981, 59). Similarly, Darwin remarked that "The survival or preservation of certain favoured words in the struggle for existence is natural selection" (Darwin 1871/1981, 61).

It is in this spirit of applying evolutionary thinking outside of biology that Helen De Cruz's Darwinian approach to mathematics explains the development of mathematical ideas by appealing to the effect of cultural selection, variation, and inheritance on the output of innate cognitive systems like the ANS. De Cruz expresses our gap problem as one of discontinuity between the properties of culturally evolved number concepts and those of our innate cognitive machinery:

> How could it be that the cultural construction of mathematics is reliable and relatively uninfluenced by circumstances, while the cognitive abilities on which it builds are imprecise and noisy? To explain these differences, we need a theoretical framework that clarifies the relationship between cultural and evolved modes of mathematical thought. (De Cruz 2007, 206)

That framework is a Darwinist approach to cultural evolution.

For De Cruz, adopting a Darwinian approach to the evolution of mathematics can help us understand how cultural evolution could have molded our genetically-inherited cognitive abilities into proper mathematical concepts: "mathematical concepts arise in human cultures through the interaction of human minds" (De Cruz 2007, 2). The idea here is that we can explain the emergence of mathematical concepts – including numbers, of course – by looking at Darwinian mechanisms of cultural evolution:

> [complex mathematical concepts] are the result of cultural evolution—a gradual and long accumulation of knowledge and mathematical skills within specific cultures… mathematical theory arises through an interplay of two systems: the cognitive system (the human brain) that is capable of entertaining it, and culture, a gradually accumulating body of knowledge to which people have contributed during many generations. (De Cruz 2007, 205)

The important question here is what the nature of this interplay is. To understand this interplay, we must grasp how De Cruz applies Darwinism to the evolution of mathematics, since Darwinism offers a framework capable of describing interactions between individuals and their cultural environments.

De Cruz thinks that adopting a Darwinian framework will allow us to understand the relationship between our innate representational repertoire and our culturally-developed practices. Darwinism allows De Cruz to frame mathematics as a product of cultural evolution, and it also allows us to consider some of our evolved cognitive modules as products of biological evolution. Darwinism thus allows us to link cognitive traits with cultural evolution. The idea is then to see how cultural evolution could have modified our innate cognitive modules to allow us to develop number concepts.

So, how can this framework help us link species-level cultural processes with individual brains? To answer this, we must first realize that, despite the fact that no individual could ever reinvent all of culture, culture is nevertheless the result of successive *individual* efforts:

> Cultural concepts are clearly more complex than evolved intuitive concepts. Yet, culture exists as a product of individual mental representations. It is the

> sum of all individual representations of a particular concept that together constitute the cultural representation of such a concept. This position is the only tenable perspective on culture, because it is the only one that allows us to examine culture through scientific methods. (De Cruz 2007, 207)

In other words, De Cruz considers culture to be a set of representations in individual brains.[136] However, given that culture is shared and communicated, and that we cannot share representations from one brain to the other in a precise and replicable manner, we need some way of sharing our internal representations with others. For De Cruz, this means accepting material expressions of inner representations as essential parts of the material incarnation of culture:

> From a materialist point of view, therefore, cultural representations only exist as stored information in individual brains (e.g., the meanings related to Guernica by Picasso are stored in our brains), and as public means of communication, which in turn have a material substrate (e.g., artefacts, patterns of sound waves (speech), public performance (drama) or writing). (De Cruz 2007, 207)

To understand the interplay between culture and individuals, we must keep in mind that neural resources are limited. Over evolutionary time, there is competition over who gets to own a piece of our neural real-estate. Our brain cannot simply engulf every cultural representation that comes its way. Rather, those cultural representations that accomplish more using less resources will have greater adaptive fit for those who possess them, since they free up additional cognitive resources for other tasks, which increased their chances of getting a lease in our heads. In short, some cultural variants are a better cognitive fit than others, as seen in the discussion of Dehaene's neuronal recycling hypothesis (section 3.3.2). One of the reasons some variants are fitter than

---

[130]Similarly, Richerson & Boyd's (2005) definition of culture also focuses on the individual: "Culture is information capable of affecting individuals' behavior that they acquire from other members of their species through teaching, imitation, and other forms of social transmission." (Richerson & Boyd 2005, 5)

others is the presence of the original owners of our brain's real-estate: our evolutionarily inherited conceptual modules, like the ANS. The fact that we approach the world through the lens offered by these modules means that some cultural variants will fare better when trying to get in our heads. Those variants that are easy to learn or that can be processed by innate cognitive modules will have greater comparative fit than those who are too complex or have no psychological appeal.

6.2.2 Darwinism and the evolution of mathematics

According to De Cruz, the evolution of mathematics has been shaped by our innate cognitive machinery in this manner: evolutionarily-inherited cognitive modules have provided content biases[137] that have put selective pressures on potential mathematical concepts, whose reception and spread depends on the degree to which they are easily represented by cognitive modules. To illustrate this, De Cruz examines the transmission of concepts of integers, negative numbers, and zero, and explains their relative cultural salience in terms of the way they fit with innate cognitive modules – more specifically, the ANS: "To explain patterns of cultural transmission of mathematical concepts, we need to examine how they interact with intuitions provided by the number module" (De Cruz 2007, 212-3). The general idea here is that those cultural practices that provide a good fit for the content of our innate cognitive modules will be easier to acquire and store, which in turn increases their chances of spreading, while those concepts that rub these modules the wrong way will have more difficulty spreading, due to the complexity of the cultural institutions needed for their transmission.

---

[137] When a cultural variant's selective fitness is explained purely as a result of how it 'fits' with our evolved cognitive architecture, this is described as the result of 'content bias', due to the fact that the bias towards selecting the behavior has to do with characteristics of that behavior.

Take the natural numbers first. Of all the number concepts, these are easily the most widespread. No culture has developed mathematical practices that has not developed natural numbers. De Cruz interprets their parallel emergence in many cultures as a sign that they can be integrated with the representations of the ANS with relative ease:

> The multiple and frequent cultural invention of positive integers can be explained as a result of their close fit with intuitions provided by the number module...The anatomical proximity of this module to other modules involved in the counting by sequential tagging procedure further adds to their salience. By being rooted in more than one conceptual module, including body part recognition and linguistic skills, positive integers are easy to learn and to transmit. (De Cruz 2007, 227)

Importantly, the modifications to the ANS brought upon by counting routines respect the format and content of approximate numerical representations, since they are compatible with a mental number line dedicated to quantity information that could be detected in the environment. In other words, even though they are more precise than the representations produced by the ANS, representations of natural numbers do not clash significantly with its output, especially when compared with other types of numerical representations, like the negative numbers, or zero.

The ANS evolved to help us keep track of the number of objects we perceive in our environment. It did not evolve to help us identify when there are no objects around us. This means that the ANS is not as well equipped to deal with zero *qua* absence of quantities as it is to yield fuzzy representations of the number of objects to which we are attending. Given that zero is not a positive quantity, is does not lie within the range of stimuli that are processed effortlessly by the ANS. According to De Cruz, this means that zero was relatively counterintuitive, and therefore took longer to develop as a numerical concept.

Going further away from our evolved representational systems are the negative numbers. While there is a sense in which negative numbers appear easy to process, since they mirror many of the natural numbers' properties, De Cruz claims that these "provide many violations which ultimately undermine our intuitions about number" (De Cruz 2007, 222). The most obvious such violation is the fact that while it is possible for us to experience positive quantities in some sense – say, by seeing four objects – De Cruz claims we cannot experience negative numerosities. In this sense, the ANS is not equipped to represent the content of negative numbers. It was only after hundreds of years that negative numbers were considered legitimate mathematical objects.[138] This is considerably longer than the time it took for zero to become a legitimate, yet peculiar, mathematical object.

In sum, reviewing historical data concerning the speed at which mathematical innovations spread shows that some practices take more time than others to spread throughout cultures, which De Cruz interprets as a consequence of the degree to which a cultural variant fits with our innate cognitive machinery. De Cruz illustrates this well with the different levels of intuitive fit provided by the natural numbers, zero, and the negative numbers. However, while it is true that variants with better fit will have a tendency to spread over populations, all other things being equal, there are many factors that can influence the extent to which a cultural variant will spread that have little to do with its fit with whether the content it expresses makes for a good fit with our innate cognitive systems.[139]

For example, De Cruz speculates that a possible explanation of the decline in progress of Chinese algebra following the Qing dynasty (1644-1911) is that its development was

---

[138]Greer, 2004, cited in De Cruz 2007, 224.

[139] As Henrich et al. (2008) argue, contrary to a view they ascribe to Sperber (1996), there are many factors that can influence the cultural fitness of a practice that have little to do with the actual characteristics of the behavior that makes up the practice. For example, whether individuals adopt a trait can vary depending on the popularity of the person they are learning the trait from, and how common this trait is in the individual's population. In this latter case of frequency-dependent bias, a conformist bias will tend to favor copying the frequent behavior, while a non-conformist bias will force the opposite behavior. See Richerson & Boyd 2005 for illuminating examples of a variety transmission biases.

hampered by a culturally-based reverence of ancient texts, which was accompanied by a reluctance to deviate from common and revered practices. The comparatively faster spread and development of algebra in other cultures that did not practice such reverence for established practices can be seen as an example where the characteristics of the practice and its fit with our cognitive architecture is not the only factor determining its cultural fitness.

The explanation of variations in cultural fitness of mathematical concepts in terms of degrees of support from evolved cognitive modules seems to be well supported by the comparative spread of natural numbers, zero, and negative numbers. Thus, with her Darwinian framework, De Cruz has a plausible description of how mathematical ideas are selected and go on to spread throughout a population, based on their interaction with cognitive biases imposed by innate cognitive structures.

6.2.3 Extended cognition and cultural evolution: where's the variation?

We have just seen how evolved modules like the ANS can determine the ease with which culturally transmitted representations like number concepts can spread through populations. Granting this, if some mathematical concepts are non-intuitive, it might appear mysterious how they can ever end up spreading through populations at all, given their lack of intuitive fit with our evolved cognitive modules. According to De Cruz, one possible explanation for the eventual spread of concepts like the negative numbers is that complex social structures allow the development of symbol systems capable of storing information outside individual heads:

> In some cultural domains of expertise, this overlap [between our brain and the content of a cultural variant] is at times so marginal that cultural transmission can only take place within a highly institutionalized context, characterized by active externalism and dedicated highly trained personnel. Without these, humans would perhaps only be able to transmit intuitive and minimally

> counterintuitive concepts. (De Cruz 2007, 228)

The idea here is that while attempting to learn non-intuitive concepts on their own may prove too costly, given the lack of fit between these and modules like the ANS, the presence of mathematical institutions, including trained professionals that communicate with culturally evolved symbol systems could facilitate this integration by offloading some of the processing into the environment. According to De Cruz, those concepts that are less intuitive will require more institutionalized contexts in order to spread, given that they benefit most from cognitive tools that reduce processing load. For De Cruz, then, adopting an extended approach to cognition is key to explaining how less intuitive mathematical ideas can spread through populations, since it is only by adopting this externalism and the associated benefits of cognitive offloading in our environment that De Cruz's Darwinian approach to mathematics explains how we overcome the lack of intuitive fit of increasingly complex mathematical ideas.

It is certainly plausible that we can explain the differences in *transmission* and *spread* of mathematical ideas by appealing to the cognitive benefits brought on by external artefacts and symbols, due to the fact that external supports allow us to perform some cognitive tasks much more efficiently. However, we have not yet seen any evidence that De Cruz' Darwinian approach can help us understand where cultural *variation* comes from. While the previous section has shown how the spread of mathematical concepts can be influenced by the degree to which cultural variants fit with evolved cognitive modules, there was no mention of the reasons or mechanisms responsible for the emergence of these variants.

Instead, the origin of those representations that manage to invade our heads via the benefits of extended cognition has been black-boxed. While uses of external objects to store previously acquired or developed information in cognition constitute legitimate applications of extended cognition to mathematical practice and the evolution of mathematical ideas, so far, it doesn't look like TXM explains the *origin* of novel mathematical content. If this is true, then appealing to extended cognition to help explain why some cultural variants manage to spread despite lack of fit with evolved cognitive modules will not help us identify how these variants are generated. For example, while De Cruz's Darwinian framework might allow us to explain why negative numbers could only spread in a symbol-enriched cultural context, it does not explain how these could emerge in the first place. In other words, so far, cultural evolution has only allowed us to explain how the presence of external cognitive support can ease cultural *transmission*, but has nothing to say about mechanisms of cultural *variation*.

This should not come as much of a surprise, given that, for the most part, theories of cultural evolution focus on mechanisms of transmission and inheritance at the *population* level (Richerson & Boyd, 2005), often neglecting the mental states of the individual (Kirkpatrick, 2009; Sperber, 2006) and the mechanisms responsible for the generation of novel content (Charbonneau, 2015, 2016). What theories of cultural evolution strive for is an explanation of why some practices spread throughout populations while others wither and die (Mesoudi 2011). And yet, without reference to an individual-level psychological process in charge of generating variation, it is difficult to see how a purely population-level description of the evolution and spread of a practice could explain the specific details underlying the emergence and cumulative change of numerical content. Much like genetic change in a species is explained in terms of genetic mutations in individuals, mechanisms of cultural evolution rely on individual-level psychological processes in their explanation of where cultural innovation and change come from. One person innovates, and, if others understand and value the innovation, it can spread via various mechanisms of cultural transmission. The innovation itself, however, originates at the individual, psychological level. So if we want to understand how numerical cognition evolved over generations, we must first understand how it could have arisen through psychological processes in individuals. For this, population-level mechanisms of cultural evolution do not seem to fit the bill.

An externalist could try and explain the evolution of numerical content by appealing to transmission or imitation errors, or perhaps to cultural mutation, but this would not change the fact that these would have taken place in someone's head—initially, at least, in a numeral-free environment—and that any change in content must be explained at the psychological, individual level rather than in reference to mechanisms of cultural evolution. To see why this is true, it is important to note that while some innovations can be the result of imitation error, such modification by error does not seem to apply to the spread of numerical content. This is because there is good reason to think that the spread of conceptual content occurs via internal reconstruction, not mere imitation:

> While the propagation of word sound may be seen as based on copying, that of word meaning cannot: it is re-productive, in the sense that it necessarily involves the triggering of constructive processes. (Claidière et al. 2014, 3).

Similarly, De Cruz, referring to Sperber (1985), writes:

> Cultural transmission, like all types of communication, is under-determined: much information is left unspecified and therefore requires a reconstruction in the mind of the recipient...each instance of cultural transmission requires a reconstruction of the concept in the individual brain (De Cruz 2011, 210)

The bottom line here is that cultural evolution begins with modification of mental states in individuals' heads (Charbonneau 2015), which, if it is understood and valued by other individuals, can then spread via imitation, re-construction and cultural inheritance. If this is true, then it is difficult to see how any modification to the content of our innate cognitive machinery can be explained by a cultural, population-level process, rather than a cognitive process at the individual psychological level.

Given that cultural evolution relies on cognitive processes at the individual level for the generation of novel content, it seems that our attempt to save externalism via this route have forced us back inside our heads. These considerations show that an appeal to cultural evolution cannot help motivate the externalist's two-stories answer to the origin of numerical content, since it cannot explain the innovation responsible for the emergence of numerical content in a world without numerical symbols, nor its spread via re-construction at the individual level, thus leaving out important details about the origin of numerical cognition.

In biological evolution, we can attribute variation to the effects of genetic mutation, which can occur due to a variety of causes, including replication errors or environmental effects. Increases and decreases to an organism's adaptive fitness can be explained in terms of the effect of such randomly occurring variations. But in the case of mathematics, it seems false to claim that variations occur as the result of randomly occurring copying errors:[140] "Clearly, in mathematics as in other complex culturally transmitted skills, variation is not generated through copying-errors, as is the case with genes" (De Cruz 2007, 267). While there is often an element of accident in mathematical discovery, in that the origins of the insight may remain inaccessible to the mathematician who innovates, it is important to realize that much of it is the result of careful, deliberate analysis on the part of the mathematical community.

More importantly, even if the construction of variants of mathematical practices were purely accidental, it would still be an individual-level phenomenon, just like genetic mutation happens in individuals before it can spread through populations. If this is true, then in order to use cultural evolution to explain how we bridge the gap, we would need an account of how individuals generate variation of numeration practices that involve representations of precise quantities and how we build on these to develop representations of natural numbers.

My claim here is that the externalist attempt to bridge the gap by providing an account

---

[140] I briefly come back to this issue of whether or not cultural variation is random below.

of the development of representations for exact quantities based on mechanisms of cultural evolution fails because cultural evolution is a population-level process, whereas generating cultural variation occurs at the level of the individual. However, a solution for the externalist here is to deny that innovation is an individual-level process, perhaps by arguing that individuals are coupled to their cultural environments, in which case cognition can be described as a population-level process. If this is an option, then perhaps the fact that there is a population-level account of cognition means we can have a population-level account of innovation, and my worry would be unfounded. To determine whether this can help the externalist, the next few sections take a closer look at Richard Menary's (2007, 2015a) brand of externalism and how it tackles the gap problem.

6.3 What's new: Innovation and enculturation of arithmetical practices

6.3.1 Innovation and enculturation

Stretching the mind further away from the head than many variants of 4E Cognition,[141] Richard Menary's (2007) brand of externalism, which he dubs *Cognitive Integration* (CI hereafter), factors in the transformative effects of culture on our plastic brains in its characterization of cognitive systems:

> cognitive integration should be understood as a thesis about the enculturation of human cognition. It is a thesis about how phylogenetically earlier forms of cognition are built upon by more recent cultural innovations (e. g., systems of symbolic representation). (Menary 2015b, 3)

---

[141] A label used to refer to approaches to cognition that share a rejection of most aspects of cognitivism and its emphasis on the inside of our head to explain how cognition works. The 4E here stands for Embodied, Embedded, Extended, and Enactive cognition. See Menary 2010b.

Reflecting the transformative effect of cultural practices on cognition, CI's taxonomy of cognitive systems includes personal and sub-personal embodied systems as well as cognitive systems that extend beyond the body, but also systems made up of the individual organism and the experts, teachers, and other aspects of the cultural niche[142] in which we evolve. Like spiders have evolved to the presence of webs, or beavers to dams, humans have adapted to an increasingly complex cultural environment, infused with practices that enrich their cognitive repertoire. On this version of active externalism, culture is part of cognition.

In the remainder of this chapter, I argue that while Menary's enculturated framework can help explain what makes the difference between numerate and anumerate cultures, it cannot help specify what makes the difference between numerate and anumerate individuals. I argue that enculturation does not have an account of innovation capable of explaining how individuals manage to improve and modify the practices of their cultural niche. This is because enculturation focuses mostly on the inheritance and transmission of practices, not on their origins, which involve individual-level understanding, rather than population-level practices and pressures. The upshot is that culture provides the necessary background conditions against which individuals can innovate. This role is crucial in the development of numerical abilities – crucial, but

---

[142] The notion of niche here is taken from a growing literature that assigns an important role to an individual organism's ability to modify its environment in charting its evolutionary trajectory, to the extent that its actions modify selective pressures on its descendants. Niche construction enters organisms into feedback cycles with their environment that alter the adaptive landscape for their (and other) species. Classical examples of niches are spider webs and beaver dams, and the surprisingly-often-discussed houses of caddys fly larvae. The idea of niche construction goes back to Lewontin (1983), who Laland et al. (2014) describe as the father figure of niche construction. For humans, an especially important case of niche construction is culture: the fact that humans are able to learn from their environment and transmit what they have learned and created – including artefacts that outlast the individuals that have crafted them – means that there are selective pressures stemming from the cultural niche constructed by our ancestors. Those who adapt to their cultural surroundings will have an advantage over those who do not. A well-known example of how cultural practices impose specific selective pressures is that of dairy farming, which resulted in increased lactose tolerance in cultures where this practice was common. See Laland et al. 2014; Wheeler & Clark 2008; Day et al. 2003.

explanatorily limited.

The rest of this chapter is divided as follows. In the next section I summarize the main lines of Menary's culture-oriented answer to the gap problem. I then discuss potential difficulties faced by Menary's enculturated approach to cognition in explaining how the first individuals that came up with numeration practices, highlighting the limitations of cultural pressures and cognitive niches in explaining individual-level differences in numerical abilities in section 6.5, while section 6.6 explores the possibility of solving these problems by framing innovation as a cultural-level process. I then highlight problems raised by Menary's culture-based approach to practices in relation to their origins and to the importance of understanding in innovation in section 6.7 before closing with a few remarks on the explanatory limits of Menary's approach.

6.4 Culture in cognition

Unlike other variants of externalism, Menary not only allows cognition to be constituted by loops that include the individual and her cognitive niche. He also allows cognition to be *controlled* by extracranial factors:

> CI has a unique position on the 4E landscape, because it is the first framework to propose that the co-ordination dynamics of integrated cognitive systems are jointly orchestrated by biological and cultural functions. (Menary 2015a, 3)

This means that some cognitive tasks can be driven by our cultural environment. According to Menary, this marks a departure from brain-centered accounts of active externalism, where the inclusion of objects outside our heads in cognitive loops with the world does not take away the central controlling role of the brain in cognition. There are many ways to go about applying externalism's cranial chauvinism to the development of numerical cognition. If all of them rely on throwing out the internalist baby with the brain-bound bathwater, in Menary's case, we are throwing out the baby, the bathwater and, to paraphrase Fodor, large parts of lower Manhattan.

This broad scope plays an important part in Menary's account of where we get our enhanced cognitive abilities from. The idea here is that the fact that our cognitive regimes have evolved to rely on our cultural niche to complete cognitive tasks can explain how we manage to develop skills like reading and arithmetic. According to Menary's enculturated view of cognition, these abilities crucially depend on the transformative effects of the practices that make up our cognitive niche. Cultural practices shape cognitive practices, on this account. They drive the dynamic integration of body and environment:

> practices govern how we deploy tools, writing systems, number systems, and other kinds of representational systems to complete cognitive tasks. These are not simply static vehicles that have contents; they are active components embedded in dynamical patterns of cultural practice. (Menary 2015a, 4)

It is easy enough to see how Menary wants to apply talk of the transformative effects of enculturation to the development of numerical cognition. While other varieties of externalism rely on numeration systems and numerical artefacts to explain how we bridge the gap, Menary's externalism takes this line of thinking further and includes the cultural practices and teachers that make up our cognitive niche as constitutive components of numerical cognition.

The idea is that the artefacts, technologies, cultural practices, and experts that make up our cognitive niche transform our evolutionarily-inherited cognitive machinery – including systems like the ANS – thus allowing us to develop a Discrete Number System (DNS hereafter) like the Indo-Arabic numerals. On this way of seeing things, our entangled reciprocal interaction with a niche populated by experts and norm-governed public practices explains how we bridge the gap because learning from our niche has transformative effects on our brain. Enculturation describes how the individual brain's integration with its cultural surroundings allows it to move beyond its biological limitations. According to Menary, this means that enculturation can answer our gap problem:

> One of the puzzles is how it is possible to move from an inherited approximate system to an acquired exact system. The process of enculturation provides the mechanisms by which such a move takes place, from the ancient capacity for numerosity to *development in a socio-cultural niche*, and *the orchestrating role of practices* in the assembly of the cognitive systems responsible for mathematical cognition. (Menary 2015,11, emphasis mine)

This certainly describes some aspects of how most people acquire arithmetical practices in numeral-enriched environments like ours: most people do learn what natural numbers are via the sustained influence of teachers that transmit practices that are part of their shared cognitive niche. However, in the next section I want to argue that there is a sense in which the fact that we learn what numbers are from our cognitive niche doesn't explain how we manage to bridge the gap. This is because, given the account of explanation in extended cognition outlined in section 5.2.3, explaining how we bridge the gap means being able to tell what the difference is between cases where the ANS is transformed and cases in which it is not. For this, appealing to the transformative influence of our cognitive niche may not be as helpful as we would like it to be.

6.5 Enculturation and the origins of the first numerical practices

If enculturation can bridge the gap, then it should be possible to explain the development of novel proto-numerical content by appealing to the effects of development in a cognitive niche enriched with cultural practices in a way that allows us to identify differences between numerate and anumerate individuals – that is, between cases of rudimentary numerical cognition and cases of proto-numerical cognition. In this section, I offer reasons to doubt that this avenue will yield the requisite difference-maker to do this. First I want to argue that there are cases where the gap is bridged despite the absence of the required niche. I then take a look at the limits of pressures and niches in driving innovation.

6.5.1 Numbers in non-numerical niches

Perhaps the most important case of bridging the gap in the absence of any cultural support is the initial innovation that led to the first proto-numerical content, mentioned in my origins problem. While any story of where and how the first arithmetical practices came to be is bound to be incomplete, there are nevertheless a few accounts of the conditions in which the first numerically-relevant artefacts capable of outlasting the individuals and cultures in which they were created came to be (e.g. Overmann 2015; Morley & Renfrew 2010; Ifrah 1998; Malafouris 2010; Everett 2017). Thankfully, there is no need here to delve deeply into this territory, since the point I am making here only relies on a what might appear to be an uncontroversial statement, namely, that someone had to come up with the first proto-numerical practices. There was a world without people thinking about proto-numbers, and then, at some point, there was a world with people thinking about and with proto-numbers. As Everett put it:

> What we do know is that someone, somewhere, at a particular moment in history, was the first person to abstractly recognize the concept of exactly five. Yet this recognition, crucial to the invention of number systems, no doubt occurred many other times independently and in various cultural lineages. (Everett 2017, 247)[143]

What is important to note here is the fact that it is possible – even, historically

[143] Similarly, Ifrah talks of "that mind-boggling moment when someone first came up with the idea of counting" (Ifrah 1998, x).

necessary, one might say – that some individuals managed to bridge the gap despite the fact that their cognitive niches had neither experts, nor public arithmetical practices, nor ready-made numeration systems to learn from.

Of course, this is not to say that there was a single individual that came up with words for the first ten numbers, for example, all at once. On the contrary, linguistic evidence tells us that there are many languages that developed words for the first few numbers but stopped around the subitizing limit.[144] At first glance, it might seem like this leaves open the possibility that proto-numerical cognition could have emerged as a result of a prolonged multi-generational process of gradual innovation, without any single individual having to come up with proto-numerical practices on their own, since it could all be a matter of incremental steps involving larger and larger names for small numerosities.

However, such an interpretation would go against the well-known discontinuity at the edge of the subitizing range that children must go through when learning the meaning of number words, discussed in chapter 4.[145] If children must labor through an Induction around the subitizing range even in a numeral-enriched world, there is reason to believe that a similar gap must have been bridged by adults living in societies where only labels for the first few numbers were available. Further, explaining the origins of proto-numerical cognition via a gradual, multi-generational process would also not explain why so many languages display discontinuities in their numerical lexicon at the subitizing range (Hurford 1987). Such discontinuities suggest that in order to go

---

[144] E.g., the Mundurucu, mentioned above in note 1. See also Hurford 1987.

[145] Recall: once children have learned to recite number words by heart, there is a prolonged learning period where the meaning of the first few number words is acquired piecemeal, in stages, for 'one', 'two', and 'three'. In this process, each stage lasts a few months. However, once they learn the meaning of 'four', children also grasp the meaning of the remaining words in their counting routine. This sudden realization is sometimes called the Induction (Rips et al. 2008b).

beyond the subitizing range and develop proto-numerical cognition, an individual must undergo some form of cognitive transformation, and that this transformation does not come easily.

It is possible that the learning process behind the Induction that allows children living in numeral-enriched environments to know the meaning of number words larger than 4 is different from the insight that allowed individuals living in numeral-impoverished environments to develop proto-numerical cognition. After all, in a sense, children learning in numeral-enriched contexts face a less onerous task, given that the environment in which their learning occurs is filled with practices and individuals encouraging them to pay attention to quantity. Granting this possibility, the point here is that whatever allowed this insight to happen could not be due to a gradual process of cultural accumulation of small increments by members of these innovative individual's niches, given the evidence showing that there is a barrier at the subitizing range that hinders the development of proto-numerical cognition, both historically and in ontogeny.[146]

If this is true, then even if we grant that the first gap-bridging occurred in cultures that

---

[146] This discussion relates to Boden's (2004) distinction between P-creativity and H-creativity, where the former describes an idea that is new to the individual, while H-creativity refers to ideas that are new to humanity as a whole. As Boden points out, H-creativity is simply a special case of P-creativity. This leads us to the question of whether two people having the same idea must share all or some of the cognitive resources required to entertain that idea. While this is a fascinating question, it need not be settled here, since the point I am making does not require the Induction that children undergo to bridge the gap be the same process that allowed the initial (H-creative) development of proto-numerical cognition. Rather, what is important for me here is that the H-creative development of proto-numerical cognition cannot have involved the numeral-and-expert-enriched cognitive niche in which the P-creative Induction takes place, since there were no numerals or experts to speak of in the H-creative case. This being said, it is worth mentioning that the presence of a discontinuity around the subitizing range in both the historical development of numeration systems and the ontogenetic development of proto-numerical cognition suggests that it is possible that the same process is responsible for bridging the gap in both the H-creative and P-creative case.

had a restricted numerical vocabulary like that of the Munduruku, individuals who managed to come up with proto-numerical content that went beyond the subitizing range did something that many cultures never managed to do, and did so without the help of a cultural niche populated by experts in proto-numeration practices. These are the individuals that managed to bridge the gap on their own, as the result of personal innovation. While precise stories of just when and where and how the first occurrences of gap-bridging behavior took place would involve more speculation than fact, it can hardly be doubted that this happened many times throughout history. In some cases, these innovations were valued and spread throughout populations, where they were improved upon by other innovators in a long process of cultural evolution that, eventually, led to the formation of formal numeration systems like the Indo-Arabic numerals.

The fact that such innovative, gap-bridging behavior occurred despite the absence of a numeral-ready cognitive niche seems to suggest that enculturation is limited in its ability to explain how we bridge the gap, since there have been cases where novel proto-numerical content can emerge in cultures that lack any proto-numerical artefacts, practices, or experts to learn from. This partially goes against Menary's claim that mathematical practices "are part of the niche that we inherit—they are part of our cultural inheritance" (2015a, 16), since such practices have not been a part of everyone's cultural niche. People nowadays only inherit mathematical practices because they were invented (and thus not inherited) by other people in the past.

There are a few ways Menary can deal with my origins problem. One is to appeal to the transformative effects of socio-cultural pressures and cognitive niches. I take a look at these possibilities in the next sub-sections. Another is to deny that the development of proto-numerical abilities constitutes an example of a discrete number system (DNS), a possibility I explore a bit later.

6.5.2 The limits of pressures

Menary could counter that such historically determinant cases as those responsible for the first proto-numerical practices could perhaps be explained as the result of individuals responding to social pressures for better quantification in a world of increasing complexity of commercial trade, for example. This is often cited as the reason why precise symbols and artefacts for numbers started appearing in various parts of the globe in societies with increasingly complex needs for measurement and quantification. As Menary put it,

> The complex social and economic pressures that required tracking exchanges involving increasingly large numbers would be the kind of socio-economic pressures that produced symbolization of quantity. (Menary 2015a, 10)

Socio-cultural pressures drive innovation. As evidence for this, we could consider the many innovations that have taken place in parallel in cultures whose technological and ideological evolution were comparably advanced – or within the same culture, as illustrated by Newton and Leibniz's parallel invention of calculus. There is a sense in which cultures sometimes arrive at periods where some ideas are in the air, ripe for the innovating. For example, many disparate cultures with no possible interaction developed the practice of using knotted string to represent numbers and perform calculations at roughly the same time (Ifrah 1998). Many have been tempted to explain the parallel emergence of notation for numbers in many cultures as the result of the need for better quantification tools prompted by an increase in the complexity of commercial exchanges.

According to Menary, innovation in such practices is due to the effects of pressures associated with changing lifestyles:

> Initially, novelty results from the pressures of increasing social and economic complexity. Small roaming bands of foragers do not need to develop symbolic number systems; post-agricultural Neolithic societies settled in villages and

towns do. (Menary 2015a, 15)

However, while such commercial pressures for better quantification could help orient innovation, this fact alone seems to fail to explain what makes the difference between those members of these cultures that have managed to bridge the gap and those who have not. The problem is that since pressures are at the cultural level, they have no way of singling out the individuals that respond to them by innovating, nor can they specify mechanisms responsible for such innovations. Of course, not every member of a community feels the pressure to innovate in the same way, nor are all innovations the result of the same degree of individual ability. For example, only a few select individuals living around the same time as Newton and Leibniz could have felt the pressures to come up with a good way of measuring the area under a curve. While the degree of technical and creative virtuosity required to bridge the gap in a numeral-impoverished environment is doubtless of a different order than that required to come up with calculus, the point here is that both innovations highlight a problem for attributing explanatory credit to social pressures, in that such individual differences between innovators and non-innovators, as well as potential differences between degrees of virtuosity required for innovations, are not captured by the effects of social pressures. So the absence of specificity of cultural pressures is a problem for Menary, since they are meant to be one of the ways in which enculturation can account for the generation of novel content:

> The DNS did not spring sui generis into the world. It did so because of a heady mixture of socio-cultural pressures, phenotypic and neural plasticity, social learning strategies, and cultural inheritance. These are the conditions for the scaffolding of the ANS, transforming our basic biological capacities into the DNS. (Menary 2015a, 15)

In this sense, the cultural niche looks like it plays the essential but explanatorily limited role of background conditions, analogous to the right chemical background that allows a particular gene to express itself a certain way, as discussed above.

However, pressures can account for differential spread of ideas through a population because they are population-level. We can appeal to the presence of such pressures in one culture but not another as part of our explanation of what makes the difference between cultures. For example, we can explain the fact that anumerate cultures like the Pirahã have not developed arithmetical practices because their culture does not value precise quantification,[147] and so there are no pressures for better quantificational practices. Similarly, the absence of such pressures, or their removal, can explain why certain practices can die off in some cultures and yet thrive in others.[148] This means that appealing to pressures for better quantification could explain difference in numerical abilities between cultures, thus answering Q2 above (section 5.2.3).

But the modification of the ANS (or other systems) happens at the individual level. The ANS is in human (and many other animals') heads. If this is true, then there is a sense in which appealing to the pressures imposed by our cognitive niche to explain how we bridge the gap cannot help identify what sets numerate individuals apart from anumerate ones. So while it is likely that these innovators developed new representational content and precise quantification tools from the effects of cultural pressures for better quantification, as Menary and others have suggested, this fact alone does not seem to allow us to explain what makes the difference between these innovators and those that did not respond to these pressures in the same way. Importantly, this specificity problem generalizes to every step that makes up the incremental, gradual historical development of the DNS, as I discuss below. It also

---

[147] For example, Gelman & Butterworth speculate that anumerate cultures do not develop words for precise quantities because "numbers are not culturally important and receive little attention in everyday life" (Gelman & Butterworth 2005, 9) in such cultures.

[148] I once again exploit the relative humility of the footnote to speculate that the disappearance of the Yuki octonary system as a result of the serial founder effect could be an instance of lack of pressures leading to disappearance of a numerical practice: the fact that Yuki culture spread out meant that changes in their lifestyles could go against seeing precise quantification as an important practice. This, combined with the decrease in numbers of experts capable of sharing their knowledge of this practice, could have led to the disappearance of this practice. See Overmann 2015; Foster 1944.

seems to apply to any intra-cultural differences including how children come to learn the DNS in ontogeny.

If this is true, then there is reason to believe that solutions to the gap problem may best be explained by accounts of individual-level processes. As mentioned earlier, there are plenty of externalist ways of doing this, most of which rely on previously-constructed artefacts present in a culturally-scaffolded cognitive niche (e.g. Dehaene 2011; Carey 2009). In such accounts the difference is one that takes place at the level of the individual. For example, Dehaene and Cohen's (2007) account of cultural recycling relies on the effects of cultural pressures, but their explanation of what makes the difference between a case of culturally-recycled individual brains and non-recycled ones is fleshed out in neural terms. Similarly, effects like the SNARC (Wood 2008) and number-related disorders like those caused by lesions (Butterworth 1999) are explainable by features of the individuals and how parts of their brains work (or don't work). The difference maker in dyscalculia and other such disorders has nothing to do with cultural pressures and can entirely be explained in reference to the individual. If this is true, then why should we expect to explain other differences between individuals who bridge the gap and those who don't by referring to anything related to development in a socio-cultural niche?

Another problem with this pressures-based account of innovation is that it seems to have the implication that we would have developed many things by now if innovation were to be explained by pressures. If pressures are the sort of thing that can be considered responsible for innovations, then it looks like we can expect things like water-condensers, cheap emissions-free cars and x-ray vision, and even reliable printers to be forthcoming given the pressures demanding their existence.

6.5.3 Niches are only as strong as their inhabitants

It seems uncontroversial that the cumulative effect of niche construction on the scaffolding of our cognitive system explains how we inherit and transmit content that emerges via innovation by other individuals. However, our interest in the gap problem concerns the *construction* of novel content that is incommensurate with its building blocks, not its transmission. There is a sense in which we shouldn't expect our cognitive niche to be able to explain the generation of novel content, given that niches can be thought of as (extremely generous) communal fridges for cognitive and cultural recipes: we can help ourselves to as much as we want of whatever we find in the fridge, but this doesn't mean that we should expect it to invent new dishes nor boil pasta for us. Menary seems to agree with this, to an extent, given that much of what he says about niches concern transmission and storage of practices, not their construction. For example, he describes niche construction as follows:

> we need a model that explains how innovations in our cultural niche are *inherited* and *propagated*, leading to changes in behavior over time. The niche construction model explains how both of these causal factors could come into play. (Menary 2015a, 7, emphasis mine)

Similarly, he describes culture as a "repository of representational systems that is passed on to later generations via learning and development" (Menary 2007,104). It is certainly not my intention to question "the importance of the environment in enhancing and supporting and amplifying cognitive capacities" (Sterelny 2010, 465). The point here is that to find the origins of (proto-) numerical cognition, we need to look for processes that generate novel content, not how it is stored and transmitted. Niches don't generate innovations, they store them and allow them to spread. Nor do niches have the potential to explain the difference between two individuals, since these differences are, by definition, at the individual level.

Menary could be tempted to reply that the focus on innovation isn't warranted. After all, virtually everyone learns arithmetical practices from others, so the focus should better be placed on learning and how knowledge is transmitted. However, this need not change the fact that the difference-maker is at the level of the individual. First, there is good reason to believe that learning conceptual content like the one associated with the DNS involves re-construction of the original content in a new individual (Claidière et al. 2014), which would suggest that many of the same constructive processes would be shared.[149] But even if this were not the case, and we found reason to believe that an innovator's path to the DNS can involve recruiting radically different processes from the learner's, the process would nevertheless still be an individual-level one, which considerations concerning niches seem ill-equipped to deal with.

These considerations based on individual differences in proto-numerical abilities and on the creative power of niches over and above the individuals that populate them suggest that foregrounding cultural aspects of cognition in trying to explain how we bridge the gap takes the focus away from the difference-maker, focusing instead on background conditions in which the gap can be bridged. This means that attempting to explain how we bridge the gap by appealing to features of enculturated cognition will be of limited use in answering Q1 above, since this approach fails to specify individual-level differences in response to cultural-level pressures.

6.6 Extending innovation through time

I have just argued that social pressures aren't specific enough to identify a difference maker between innovators and non-innovators at the individual level, since the same

[149] Similarly, De Cruz, referring to Sperber (1985), writes: "Cultural transmission, like all types of communication, is under-determined: much information is left unspecified and therefore requires a reconstruction in the mind of the recipient...each instance of cultural transmission requires a reconstruction of the concept in the individual brain" (De Cruz 2007, 210).

pressures can affect individuals in radically different ways. However, Menary can perhaps avoid this problem by framing innovation as happening at the social level.

This could then allow Menary to accommodate the apparent exception provided by the development of proto-numerical practices in a non-numerical niche, since such innovations in a numeral-free world could fail to be considered as proper examples of a DNS, which, it could be argued, results from cognitive processes that take place over much longer historical time periods. In this section, I explore this possibility and find it wanting, because seeing innovation as a social phenomenon only works by building on innovation that occurs at the individual level.

6.6.1 Is innovation social?

According to Menary, the previous section has things in reverse: contrary to my claim that individuals are responsible for innovations, he argues that innovation occurs at a social level. For example, in response to Fabry's (2015) suggestion that predictive processing (e.g. Clark 2013) could help ground the enculturated approach at the neural level, Menary writes: "the predictive processes at a sub-personal level cannot be driving the *innovations at a social level* that lead to enculturated cognitive systems" (Menary 2015b, 1, emphasis mine). Menary worries that the individual-level focus of predictive processing clashes with the explanatory emphasis of enculturated cognitive integration, which is "the population-level effects of normative patterned practices (henceforth NPP), such as mathematical practices" (Menary 2015b, 1).

According to Menary, the fact that predictive processing aims to describe online processing and explain it in terms of error minimization means that it cannot explain innovation in normative practices, since the pressures that lead to these operate at the social level. If the brain is engaged in predictive error minimization (as sub-personal processing) in the here-and-now, then it cannot be driving the innovation of new NPPs over many generations, which are found at the social, or populational, level. Given Menary's claim that cultural pressures drive innovation, the claim here seems to be that whatever drives innovation is at the same level of cognition as innovation itself, and so innovation happens at the social level.

Seeing innovation as a cultural-level process certainly applies to the majority of technologies and tools that abound in the world around us: one would be hard-pressed to find a cultural product that was the result of a single individual's work. As Hutchin's (1995) description of the historical development of navigational tools and techniques and Richerson & Boyd's (2005) many case studies beautifully illustrate, most tools, artefacts, and practices are the result of a process of historical accumulation of small improvements on previous design. Henrich et al. (2008) offer a representative quote from historian Joseph Needham who summed up the cultural origins of the steam engine thus: "No single man was the father of the steam engine; no single civilization either" (quoted in Henrich et al. 2008, 130, from Basalla 1988).[150]

The same sort of stepwise process of cumulative improvement to previous innovation applies to the historical development of the DNS. For example, Ifrah describes the development of the modern Chinese numeration system as the "fruit of a veritable cascade of inventions and innovations". He adds, "It emerged little by little, following thousands of years during which an extraordinary profusion of trials and errors, of

[150] Without claiming that the origins of numerical practices are instances of the myth of the heroic inventor (see Henrich et al. 2008 for references behind this expression), it should be noted that there are some circumstances where inventions have immediate, sudden impact on their societies. Not all innovations are created equal, and there are many cases throughout the history of mathematics of such individual-led game changers. One could consider Cantor's work on transfinite numbers an example of such innovation. In a somewhat less grandiose, more culinary context, Dominique Ansel's invention of the cronut – a cross between a croissant and a doughnut – could perhaps also count as such a game-changer. Both cases seem to illustrate how "major transitions in society need not await a series of innovations, each of small effect, but may result instead from key innovations or from coordinated flexibility in response to changing conditions" (Laland et al 2014, 12).

sudden breakthroughs and of standstills, regressions and revolutions occurred" (Ifrah 1998, 674).

This historical, extended timescale on which such cultural processes of innovation take place seems to support Menary's claim that innovation is social. If this is true, then one can indeed argue that the first times people thought of something with content like 11, or "the quantity that can be put in correspondence with both hands and a finger", they weren't in possession of anything like a modern-day DNS like the Indo-Arabic numeration system that dominates our arithmetical practices.

Of course, if there is little reason to think that the practices associated with the Indo-Arabic numerals *could* have emerged as the result of a single individual's eureka moment, there certainly is no reason to think this *is* how it came to be, since the historical records show a gradual, stepwise accumulation of small improvements stretching for many generations for the system to arrive at its current form (e.g. Ifrah 1998). So if the gap-bridging story is the historical transition starting from the ANS and ending with the DNS, then the fact that the first step in this process was the result of individual innovation does not count against the enculturated answer, since such primitive cases are not true gap bridging stories. So it looks like seeing innovation as a cultural-level process takes the bite out of the origins concerns I presented above.

However, it is important to note that the cumulative construction of numeration systems and other artefacts mentioned above is the result of scaffolded steps of *individual innovations*. So while there is a sense in which we can accommodate Menary's claim that innovation occurs at a social level, this is only possible insofar as social innovation is made up of an accumulation of individual innovation of the sort present in the origins of proto-numerical practices. If this is true, then the origins problem generalizes to every individual step in the historical development of the DNS, since each step represents an individual's small innovation, for which the enculturated approach seems incapable of specifying a difference maker, given its explicitly population-level focus.

So even if we were to deny that the first cases of precise quantificational systems lacked sufficient complexity and expressive power to count as proper DNS[151] we still need to explain what makes the difference between a person that understands what a limited, primitive version of numbers are and one who doesn't. Given that, for every step in the cumulative historical development of the Indo-Arabic DNS, only a select few individuals managed to innovate their quantificational practices despite sharing a cognitive niche with many non-innovators, the cultural setting seems to serve as a necessary backdrop for the innovation to happen, but not a difference maker. Rather, it is the individual-level process responsible for the innovation that makes the difference between a person that understands what (proto-) numbers are, and one who doesn't, at each step of the development of the DNS. Of course, this is not to say that each step in the historical development of the DNS or any other product of cumulative cultural evolution requires the same creative input. The point here is that looking at innovation as a social process of cumulative individual contributions obscures the individual-level processes behind these contributions, big and small. In a slogan: If individuals don't

[151] Presumably, this denial would require coming up with criteria for when a numeration practice is truly, well, numerical. Perhaps we could want to base ourselves on the quantities reached, or the portability, ease of use, ease of comprehension, or any combination of these and other factors – not a simple task, to say the least. Moreover, this sounds like we would end up excluding some systems and including other systems of comparable expressive power, which would make hard lines somewhat arbitrary, as I mentioned in the introduction. It would appear senseless to distinguish systems based on the number of items they have developed symbols for. For example, it seems useless to distinguish between systems that have symbols for up to 27 items, and systems that have symbols for up to 29 items, say.
On the other hand, there is a good reason to distinguish between systems whose expressive power is limited by the absence of formalized generative syntax, as I have proposed here with the term 'proto-numerical', and systems that require understanding a generative rule for their mastery, rather than rote memorization. This relates to Hurford's (1987) talk of numeral lexicons that lack the syntactic resources needed to express larger numbers. The mastery of the syntax of a formal system would seem to represent a separate step in the development of formal numerical cognition that requires separate cognitive processes, such as those associated with multi-digit numeral processing (e.g. Nuerk et al. 2015). Similarly, it may be useful to distinguish systems that cannot be extended indefinitely from those that can be extended indefinitely, as is the case for the Indo-Arabic numeration system. The reasons for such distinctions go beyond our present concerns, however. See Schlimm 2018 for work on these topics, as well as Beller & Bender 2008 and Overmann 2015, 2016.

innovate, cultures don't innovate.

6.6.2 Complementary levels of innovation?

Perhaps denying that innovation occurs at the cultural level is too harsh: why not say that innovation happens at both individual and cultural levels? Given Menary's inclusive approach to cognition, which accepts that cognition occurs at many levels, we could accept that innovation happens at many levels, and that the innovation that matters is the one that concerns cumulative historical construction like the one behind the DNS.

While this may be appealing, prioritizing the cultural development of the DNS seems to take things the wrong way round, since it is only possible given the individual-level process of innovation that makes up each step of the historical development. Moreover, and perhaps more importantly, there are reasons to believe that seeing innovation as a cultural-level process has unfortunate consequences, including how we frame the individual input in the construction of public practices.

One reason to doubt seeing innovation as primarily social is that this seems to go against common practice. While philosophers are unfortunately all-too-comfortable with radical reformulations of common notions, the fact remains that we do speak of Leibniz and Newton as the individuals that came up with calculus, rather than attributing this to their specific culture.

Seeing innovation as a public-level processes also seems to clash with the fact that innovations often happen as the result of sudden realizations on the part of individuals. Of course, there are cases of inventions that take years to take shape, and, as already noted, most artefacts and practices that make up our cognitive niche are the result of generations of tweaking. But this does not negate the fact that, in many circumstances, individuals happen upon a discovery – more often than not, by accident – and in such cases the innovation results from sudden eureka moments, often following incubation periods (Chen & Krajbich 2017; Hadamard 1945).

For example, Poincaré (1910) claimed he came up with a proof for a property of Fuchsian functions when he put his foot on a bus, after weeks of unsuccessful attempts, while Paul McCartney claims he woke up one day and the melody for the classic song *Yesterday* was there in his head. While bridging the gap does not require the same creative genius as Poincaré (or, perhaps to a lesser extent, McCartney), a similar process of realization following incubation applies to the Induction that leads children to proto-numerical cognition. Importantly, such eureka moments occur at the individual level.

Again, the important point here is that the realization that accompanies eureka moments, big and small, happens at the individual level. So even if the overwhelming majority of cultural artefacts are the result of cumulative modifications that stretch through generations, this does not mean that innovation is best seen as a population-level process, since each modification is done by individuals. Much like mutation is something that can be described at the level of individual genes even though the gene's environment can play a major causal role in the process, innovation can be described at the level of individuals even if Menary is right that cultural pressures drive these. If this is the case, then enculturation seems limited in its explanatory power with respect to how we bridge the gap, since the role of culture in innovation is to provide a backdrop against which individuals innovate. To frame innovation at a social level seems to ignore the key contribution of specific individuals in the historical accumulation of innovations.

6.7 Innovation and Public Practices

In this section, I want to explore another aspect of enculturation that seems to have negative consequences for its ability to explain innovations like the ones responsible for the construction of the DNS. More precisely, I want to argue that Menary's focus on the active role of our cultural niche in control of cognition and the associated claim that practices are cultural-level parts of cognitive systems black-boxes the essential transformative effects of individual-level innovation (and re-construction).[152] There are two related problems with seeing cognitive practices as essentially cultural, both of which have to do with the origins of such practices: the first is that seeing practices as cultural level obscures their origins, while the second is that seeing them as public obscures the role of understanding in innovation.

As mentioned above, Menary's culture-oriented take on cognition sees public practices like those involved in reading and arithmetic as having control over some cognitive systems: "In enculturated systems, the really important work is being done by the processing governed by normative patterned practices whose properties are understood primarily at the social or populational level" (Menary 2015b:8). This means practices are essentially cultural in nature: "Cognitive practices are culturally endowed (bodily) manipulations of informational structures" (Menary 2015a, 4).

The problem is that framing cognitive practices at the population-level seems to pass over the fact that it is individuals who first invent these practices before they spread

---

[152] This is not to imply that the processes responsible for an innovation are the same ones involved in the re-construction of the original content in a learner's head. While I think there is a way of reading Sperber (1985) or Claidière et al. (2014) that is consistent with this claim, the point being made here does not rely on taking a side on this issue. What must be kept in mind here is that both the innovative process and the re-constructive one are black-boxed if we construe practices as essentially cultural, as Menary seems to want to do.

throughout populations.[153] For example, consider Richerson & Boyd's (2005) description of how cultural evolution can overpower biological evolution in the case of the unpleasant taste of medicine: while our taste buds might have evolved to detect bitter, disgusting plants, thus protecting us from potentially poisonous sources of food, cultural transmission of the knowledge that some unpleasant tastes can be compensated by medicinal benefits has led to the spread of behavior that goes against our biologically-driven aversion to such tastes. While the mechanism of transmission of this practice is at the population level, this practice, like many others, starts off with a single individual's inquisitive and innovative behavior:

> We take our medicine in spite of its bitter taste, not because our sensory physiology has evolved to make it less bitter, but because the idea that it has therapeutic value has spread through the population. In the distant past, some inquisitive and observant healer discovered the curative properties of a bitter plant. Then a number of processes … might cause this belief to increase in frequency, despite its horrible taste. (Richerson & Boyd 2005, 11)

This illustrates the role of individual-level psychological processes in the generation of novel cultural practices. Menary seems amenable to the idea that individuals innovate novel practices when he writes that "the organism is predisposed to manipulate its environmental niche, or in some cases *create* it" (Menary 2007, 103; emphasis mine). But seeing them as essentially cultural seems to go against the driving role played by individuals in the innovations that, in some cases, lead to population-level patterns of behavior that Menary considers drivers of some cognitive practices. This seems to suggest that framing practices as population-level black-boxes the role of individuals

---

[153] It is controversial to claim as I do here that individuals invent practices, given that many frame practices as being constituted by structures of interactions between individuals, which would seem to prevent them from being invented by a single individual. The point I am making here does not rely on rejecting this approach to practices, since even if practices are constituted by such social interactions, they are still often framed as patterns of individual performances (e.g. Rouse 2007). Thus, it seems acceptable to claim as I do here that such patterns of individual performances have origins, and that these origins are performances by *individuals*.

in their origins.

Perhaps more importantly, taking the view that practices are population-level seems to force us into a problematic expert regress if we try to identify their origins. For example, if we want to explain the shift from the ANS to the DNS in terms of development in a niche where experts abound, we have to be able to explain how these experts came to acquire the practices that they teach others. Since it can't be experts all the way down, we need to find an explanation of the origins of this expertise. For this, I suggest, we need to appeal to individual-level innovation strategies.

Another reason to resist framing the origin of practices at the cultural level is what we could call lone wolf innovation, which occurs when an individual innovates a practice on their own and then this innovation fails to catch on. In such cases, the behavior associated with the practice is manifested, and yet the fact that it has not become a public pattern of behavior seems to bar it from being considered a practice. Similarly, consider cases where two individuals both invent the same modification to a current practice, but only one of these catches on. If the first to invent the practice fails to popularize, while a second with the same innovation does, is the less popular variant not an innovation because it isn't a crowd-pleaser?

In addition to the problems related to the origins of practices, a related problem with seeing practices as population-level patterns of action is that this seems to ignore the essential contribution of individuals and their understanding of these practices in the process of innovation: when an individual makes an improvement upon a previously learned practice, she is not acting blindly. Even in cases of accidental discoveries, the individual that happens upon a novel state of affairs is interpreting the event according to an understanding of the situation.

For example, Alexander Fleming's accidental discovery of penicillin was due to the fact that he had left some petri dishes out too long, and mold only grew well on some of them. Though he did not realize the full potential of his discovery, he did grasp enough of it to explore novel uses of mold against bacteria. This would not have happened to a person without a previous understanding of how these interact.

Similarly, when individuals come up with ways of quantifying their environment or improving upon quantificational practices they may have learned from others, they are not acting blindly, merely imitating the gestures of others. They are acting according to an understanding of a task they are trying to accomplish. This understanding seems lost when we frame practices as population-level patterns of action, since understanding does not occur at the population level. Consider Menary's description of long multiplication:

> Practices are patterns of action spread across a population. However, I am inclined to think that practices are not simply reducible to the bodily actions of individuals. Whilst doing long multiplication requires a bodily action of me, what I am doing cannot be described exclusively in terms of those bodily actions. The practice is a population, or group level phenomenon, not an individual one. (Menary 2015a, 4f)

Menary here claims descriptions of practices in terms of bodily actions must be framed within a group level phenomenon. Going in the opposite direction, I would argue that what is missing is an individual-level understanding of the rules that make up long multiplication. Bodily actions on their own are not enough to certify mastery of a practice – at least, not one like arithmetic. This is because the practice of arithmetic – even the basic, proto-numerical skills that could have emerged in the absence of a niche with proto-numerical practices – requires understanding of rules. Unlike many other culturally-inherited practices like driving cars or the Macarena, the practice of arithmetic requires more than physical manipulation and copying what others do, since the learning individual must come to an *understanding* of a recursive rule.[154]

There is a sense in which anyone can blindly copy the symbol strings and utter the sounds associated with the practice of arithmetic. But few would consider this ability as a case of mastering arithmetical practices. For example, children learn to blindly recite initially-meaningless lists of number words long before they understand what numbers are (Carey 2009). However, once they have figured out what the first few number words mean – in a slow, piecemeal process (Wynn 1992b) – they soon enough generalize to the rest of their count list. Unlike learning how to drive, which doesn't require us to know anything about how cars work, the correct practice of arithmetic and the right way to apply number words requires understanding what numbers are, which is an internal, individual happening. This is might be why learning to count takes so long and why no animal has managed to generalize addition using symbols, despite a few cases of decades-long training and rote learning of associations between symbols and quantities (e.g. Matsuzawa 1985; Pepperbeg & Gordon 2005).

If a person knows how to do basic formal arithmetic, they understand how to generate the right symbols in the right order. They have mastered rules. This is what explains that there can be productive use of practices like arithmetic: when individuals come to master arithmetical practices, they can perform instances of such practices that they have not been shown. Generalizing this, I want to argue that when they make improvements to their repertoire of cognitive practices, in at least some cases, individuals are acting out of an understanding of what needs to be improved. If this is true, then seeing practices as socially-based seems to obscure the important role of individual-level understanding that drives innovation.

Further support comes from the observation that cultural pressures drive innovation:

---

[154] See Rouse 2007 for more on the importance of understanding in the transmission of practices.

how could cultural pressures drive innovation unless individuals understood the improvements that pressures are guiding them towards? Of course pressures do drive innovation, in the sense of providing a backdrop against which individual innovation can modify current practices. But this suggests that innovation is not accidental, that it results from an understanding of what practices are available, of their limitations, and of the practices that would best satisfy such pressures. But if we see learning arithmetical practices as being 'on paper' before they get in our head, as Menary does, it looks like we aren't taking the understanding required for innovation seriously.

In summary, the requirements imposed by the important role played by realizing and understanding in different stages of the development of numerical cognition seems to block seeing numerical practices as population-level, since the understanding that leads to mastery and that led to innovation takes place at the individual level.

6.8 The explanatory limits of cultural evolution and extended cognition

It is important to note here that I am not denying the importance of enculturation for the development of bodies of knowledge like mathematics. Clearly, no single person could ever accomplish what we do as a species. But acknowledging the importance of enculturation for the development of mathematics and arithmetic does not necessarily explain it. The fact that humans gradually accumulate innovations over generations is indeed a cultural process, one that is undoubtedly responsible for the incredible achievements of mathematics, and science as a whole. But this cultural evolution relies on *individuals* responding to their (cultural) environment and building on it. It is this building process, the generation of novel content, innovation, that needs to be explained in our case. For this, appealing to the fact that human cognition extends into the environment to include cultural processes is limited in its explanatory power.

To understand why, it is important to realize that the emergence of the first number concepts took place in a world that did not contain any explicit, external symbols for numbers. And yet, there is no question that numerical content did emerge in an enculturated context, probably due to the demands of increasingly complex commercial exchanges and practices which benefitted from keeping tallies on precise quantities. But the fact that an event or practice takes place within an cultural context does not mean that enculturation can meaningfully explain its origins. Similarly, while it is possible that embodied cognition is a pre-requisite for the development of numerical content, as those who attribute an important ontogenetic and phylogenetic role to the use of fingers to explain the origin of numerical content believe, this does not help narrow down the process that led to the construction of novel numerical content discontinuous with the output of our innate cognitive systems. As mentioned above, the fact that there are people with hands and no number concepts looks like an easy way to illustrate the limitations of this approach, as does the fact that the available evidence seems to suggest that counting on one's fingers is not necessary to develop numerical abilities (Crollen et al. 2011).

If these considerations ring true, then the level at which cultural evolution and extended mind frameworks apply does not look like it is specific enough to explain how we bridged the gap. In the case of extended cognition, not only does the framework not allow us to identify the how artefacts with numerical content came to be, but the construction of novel numerical content itself appears outside of this approach's explanatory reach.

A good way to illustrate that it is not always explanatorily rewarding to appeal to enculturation or extended cognition, irrespective of whether or not they describe how our minds work, is to consider our digestive system. Citing Wrangham (2009), Sterelny (2010) points out that our jaws – and, relatedly, our brains – have evolved to their present shape because we started cooking our food a very long time ago.[155] In a sense, then, our digestive system could be described as being extended beyond our stomach, given that we essentially rely on external artefacts to process, select, ferment, and store food in ways that have profoundly affected our digestive system and jaw:

> The physiological demands on hominin jaws, teeth and guts have been transformed by cooking and more generally by food preparation and food targeting...[w]e are obligatorily cooks. Moreover, we supplement cooking by pre-engineering our food sources. (Sterelny 2010, 467)

Similarly, given the fact that cooking requires fire, and that keeping a fire going while hunting prey is not a task fit for a single individual, one could argue that the shape of our jaw and the ensuing increase in brain size that followed the jaw's gradual evolution are the results of enculturation, since such changes to our cooking practices could only have occurred in societies where the right tools and social hierarchies were developed.

And yet, despite the fact that our digestive system is entirely dependent on technology that has allowed us to pre-process our food outside our body for thousands of generations, this information is not required to explain why certain foods cause heartburn, or other facts about how our digestive system works: "there is no explanatory mileage in treating my soup pot as part of my digestive system, once its importance as a scaffold is recognised" (Sterelny 2010, 468). Nor does our essential reliance on external artefacts for our food intake mean that such technology is part of our digestive system. As Sterelny put it in this delicious passage:

> Our digestion is, then, technologically supported in profound and pervasive ways...We have engineered our gustatory niche; we have transformed both our food sources and the process of eating itself. Our under-powered jaws, short gut, small teeth and mouth fit our niche because we eat soft, rich and easily digested food. Our digestive system is environmentally scaffolded. But is my

[155] Sterelny mentions that while estimates locate the advent of cooking food anywhere between one and two million years ago, there is no doubt that we have been cooking our food for at least 400 000 years.

> soup pot, my food processor and my fine collection of choppers part of my digestive system? As far as I know, no one has defended an extended stomach hypothesis, treating routine kitchen equipment as part of an agent's digestive system; indeed "extended stomach" and "extended digestion" on Google Scholar return some very strange hits. (Sterelny 2010, 467-8)

So while we may indeed rely on external objects (and people) for our daily cognitive regime, including learning and using numerical content, this does not mean that we can explain the emergence of novel content by appealing to these general facts about how our minds work. It is not always helpful to appeal to the external aspects of our minds to explain particular phenomena, even if our minds are indeed extended, enculturated, embodied, or any other variant of active externalism.

CONCLUSION

In his review of Susan Carey's *The origin of concepts* (2009), Nicholas Shea writes:

> The philosopher of mind who stands convinced of the relevance of empirical results soon hits a problem. Lured in by a few interesting studies, the door opens on a cacophony of data, like a frenetic party in full swing. There are just so many studies out there. How do they fit together? And what do they all add up to? (Shea 2011, 129)

Shea goes on to praise Carey for her ability to make sense of all these disparate data – rightly so. In my thesis, I have tried to navigate through such a cacophony of data concerning the potential origins of numerical cognition in order to paint an internalist-friendly picture of how numerical content of the sort used in the practice of arithmetic emerges from evolutionarily-inherited systems like the ANS and the OFS. This required filtering data from a motley crew of research domains, from anthropology to neuroscience, from comparative psychology to archeology, in order to determine the ability of the dominant externalist approach to numerical cognition to weave an explanatorily satisfactory account of the origins of mathematically-viable concepts of numbers.

A quick recap will help us see the ground we have covered. In the first two chapters we saw that recent breakthroughs in a wide variety of fields have allowed us to identify innate cognitive systems capable of producing representations with quantitative content. We also saw that there is good evidence that systems like the ANS and the OFS are recruited in numerical tasks, which suggests that we may have found leads on the potential neural underpinnings of our ability to think about numbers. However, we also saw that these systems' limitations mean that they cannot, on their own, represent numbers.

This led us to take a look in chapters 3 and 4 at how prominent thinkers in the study of numerical cognition like Stanislas Dehaene and Susan Carey exploit their knowledge of the extensive literature concerning these and other evolutionarily-inherited representational systems to explain how we overcome their limitations and represent natural numbers. This allowed us in turn to see the extent to which their accounts rely on things outside our head to bridge the gap between our limited evolutionary heritage and the unbounded expressive power of arithmetically-viable representations of number. I then presented my origins problem to these accounts, arguing that taking an externalist approach to the origins of numerical cognition is tantamount to putting the cart before the horse, since it relies on the presence in the environment of things whose existence in turn relies on the very same internal representations whose origins we are trying to explain.

In chapter 5, I took a closer look at the philosophical motivations behind this externalism in Clark & Chalmer's classic paper on the extended mind before exploring the constitutivity of external supports for cognition, as framed by Catarina Dutilh Novaes. This led me to discuss the relationship between external and internal representations for numbers at both ontogenetic and historical timescales. To ground this discussion, I took a look at Lambros Malafouris' externalist account of the development of abstract notation for numbers in Sumeria as well as the limited – arguably, nonexistent – numerical abilities of cultures like the Mundurucu and the Pirahã . These case studies helped me argue that externalist approaches have trouble explaining what makes the difference between individuals that manage to bridge the gap and those who do not, because they rely on the existence of cognitive loops with things outside the head to explain this difference, while I claim that both individuals have access to the same external resources, so the difference maker must be inside their head. This led me to claim that externalism about numerical cognition leaves out important details about the origin of numerical content, especially, the concept of precise quantity, and that an initial segment of the natural numbers seems out of reach for externalist approaches to numerical cognition, given their inability to explain the difference between individuals that have bridged the gap and those who have not, in both numeral-free and numeral-enriched environments.

In the last chapter, I considered two responses to this charge. The first is to try to explain the emergence of novel numerical content by appealing to mechanisms of cultural evolution, as described by Helen De Cruz. I argued that this doesn't help, since such mechanisms are population-level, while the generation of novel content occurs at the level of the individual. Similarly, in response to Richard Menary's enculturated approach to numerical cognition, I argued that innovation matters in how we want to answer the gap problem and that enculturation is not well equipped to answer the individual-level aspect of this problem. Rather, enculturation is better suited to explain the cumulative growth of arithmetical practices and their spread through populations, as well as what makes the difference between two different cultures' numerical abilities. The claims I have relied on have to do with the fact that there are cases where we bridge the gap without help from our niche (as in the case of the original development of proto-numerical practices by creative individuals) or we can't bridge the gap despite being in a niche (as in cases of dyscalculia). I also argued that each step in the cumulative construction of cultural practices and artefacts is the result of *individual* innovation, and that the fact that the transformation occurs at the individual level means that social-level innovation and practices cannot supply a difference maker explaining what sets innovative individuals apart from non-innovators.

In my analysis of the externalist approaches of Dutilh Novaes, De Cruz, Malafouris, and Menary, I relied on claims about the relation between mental content and symbols, the continuity of ontogenetic processes through the historical development of numerical cognition, and the role of individual-level psychological processes in mechanisms of cultural evolution. Even if these claims are mistaken, these externalist frameworks could only help rid us of the apparent methodological circularity involved in taking a manifestation of numerical cognition to be one of its causes, and thus externalism would still not have a full story about the origin of numerical content in a world without numerical artefacts.

Of course, I would be foolish to deny the essential role of external symbols for numbers and public practices in shaping our arithmetical skills, given that learning from our interaction with a niche in which such practices abound is clearly behind our advanced numerical abilities. However, none of what I have argued for denies the important role of external supports for cognition in the advanced practice of arithmetic. Rather, what I have argued here does not negate that some arithmetical practices – especially the ones mathematicians care about, and the ones that deal with 'large' numbers – could be considered as cases of extended, enculturated, embodied, or enactive cognition, given that our feeble brain-bound memories are not strong enough to allow us to keep track of precise quantities beyond a certain unidentified limit.

However, if I am correct in stressing the inability of externalist accounts to answer my origins problem, while some cognitive processes – including arithmetic involving sufficiently high numbers – may very well be usefully described as being extended into the world because they are impossible to complete without external aids, the role of these artefacts could still be limited to cognitive offloading of representations that were constructed via processes that can be described in terms of purely internal resources. While in current-day arithmetic and mathematics, this role is of capital importance, given the quantity and variety of arithmetical tasks that involve large numbers on a daily basis, it should not cast a shadow over the painstakingly-acquired understanding of numerical quantity on which it rests, without which there would be no gap problem to speak of. I have tried to argue that this understanding allows the development of at least an initial segment of the natural numbers.

If I am right, then it looks like the externalist point of view cannot give us the full story of how systems like the ANS and the OFS allowed us to develop mathematically-viable number concepts. Thus, a different, internalist approach to the origin of number concepts seems like a worthy avenue to explore, if only to explain the initial, artefact-free development of numerical content that externalism leaves out.

The question is: can internalism fare any better? To show internalism is a promising framework to help us bridge the gap, it would be necessary to have an account what sort of intra-cranial processes and systems could be involved in the development of numerical content. Given the important role of attention in the externalist accounts of numerical cognition considered here, as well as in many explanations of how we interact with external objects in extended cognition loops (see e.g. Clark's 2008 discussion of language as an external resource), I will allow myself to speculate that attention to quantity is a worthy avenue to explore.

While the issue of whether it is best to explain the development of numerical content by adopting an externalist framework may not allow us to explain how numbers get the objectivity and generality that made Frege and Husserl weary of any psychological intrusion into foundational issues, settling this issue could certainly represent progress in that direction. This is because once we have a naturalistic understanding of what numbers are, it seems reasonable to hope that their generality, objectivity, and apparent mind-independence could be explained in terms of the general properties of the cognitive processes that generate them.

One could argue that such an attempt to explain the unique properties of numbers in terms of psychological processes was already present in the work of Dutch mathematician L.E. J Brouwer, who tried to build foundations of mathematics that rely not on sets and axioms, but on the recursive properties of consciousness. Regrettably, Brouwer's intuitionist approach of the construction of number by the human mind was couched in a language laden with mystical notions (van Stigt 1991), and while his description of the origin of numbers did produce mathematically viable tools – at least, as far as the natural numbers are concerned – his account did not bother with the actual details of how the mind works.[156] However, given that there are cognitive processes that share the recursive properties of Brouwer's account of construction of natural numbers from the primordial intuition of time, such as attention (Dehaene & Changeux 2011), internalists might gain from trying to adapt some aspects of his views to a more empirically-informed framework to see if it is possible to provide a psychologically-viable version of his intuitionism.

These are, of course, speculative remarks aimed at illustrating how an internalist approach could help bridge the gap without attributing a constitutive role to external artefacts. And yet, even if these suggestions turn out to be ludicrous, the problems with externalism listed above suggest that an internalist approach is warranted if we wish to describe all cases of ontogenetic development of numerical content. However, this would not mean that externalism needs to be abandoned. On the contrary, examining the causal role played by external symbols in numerical cognition can help us understand the internal processes involved in the development of number concepts, since knowledge of this causal process may help identify internal analogues that play the same role as external symbols. So while relying on external symbols might not give us the whole story, it can move us in the right direction.

In any case, as things stand, neither internalist nor externalist accounts have managed to exploit the vast amounts of empirical data to offer an answer to the gap problem. Perhaps more importantly for the *philosophical* study of the foundations of arithmetic, neither approach has delivered an empirically-informed account of what numbers are

---

[156] In Brouwer's defense, few such details were available at the beginning of the 20th century.

that could offer sufficiently-detailed explanation of where numbers get their generality and objectivity to offer an appetizing alternative to the platonist.

The first step in motivating a philosophical shift away from platonism using the vast numerical cognition literature is to show that it can provide a workable alternative to the platonist story. As Dehaene admits, "the burden of elucidation clearly falls upon psychologists and neuroscientists, who will have to figure out how a finite brain, a mere collection of nerve cells, can conceive such abstract thoughts" (Dehaene 1997/2011, x). Given the incredible progress of this domain of research despite its young age, on this front, there is room for optimism. For example, consider the fact that all the writers covered in this thesis have a commitment to Darwinism about mathematics, in that they view mathematics as the product of Darwinian processes of biological and cultural evolution. Already, this marks a welcome departure from the foundationalist program that dominated philosophy of mathematics in the last century, in that it offers at least the possibility of coming up with a naturalized account of numbers.

Darwinism looks for causes to behavior in the history of a species. In so doing, it embraces science's general commitment to causal explanation – a notable improvement over traditional foundations of mathematics. Darwinism also accepts that historical development is often a series of accidents and contingencies, and that the products of this development are fleeting and imperfect. On this point of view there is no room for a perfect platonic world proposed by Frege and Husserl.

Thus, describing mathematical thought as the product of a Darwinian process of evolution allows us to look at mathematics from a naturalistic perspective, where mathematics is considered as a growing body of knowledge produced by human activity instead of a static, platonic world of eternal truths that we somehow discover. Seen through the lens of evolutionary theory, mathematical concepts, their origins, and their evolution are all subjects of scientific inquiry. This naturalistic perspective shifts the focus of enquiry from considerations internal to mathematics that were typical of traditional approaches to the foundations of mathematics to the actual human practice of mathematics and the role of human agency in the emergence of mathematical concepts.

An advantage of Darwinism, then, is that it comes with ontological and epistemological commitments that make it possible (in theory) to explain how humans could interact with mathematical objects, a question that early 20th-century foundationalist programs, with their focus on internal consistency and foundational concepts, failed to address. The fact that this approach seems shared by philosophers reflecting on the origins of numerical cognition could indicate that we are gradually shifting away from the rampant platonism of classical foundational projects. As a naturalistically-inclined philosopher of mathematics, the fact that Darwinism even listens to the cacophony of numerical cognition data is music to my ears.

REFERENCES

Aach, J. (1990). Psychologism reconsidered: A re-evaluation of the arguments of Frege and Husserl. *Synthese* 85: 315-338

Agrillo, C., Piffer, L., Bisazza, A., Butterworth, B. (2012). Evidence for Two Numerical Systems That Are Similar in Humans and Guppies. *PLoS ONE* 7(2): e31923. https://doi.org/10.1371/journal.pone.0031923

Agrillo, C. (2015). Numerical and Arithmetic Abilities in Non-primate Species. In R. Cohen Kadosh & A. Dowker (Eds.), *The Oxford Handbook of Numerical Cognition*. Oxford: Oxford University Press.

Allport, A. (1993). Attention and Control: Have We Been Asking the Wrong Questions? A Critical Review of Twenty-five Years. In *Attention and Performance XIV (Silver Jubilee Volume): Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience,* 183–218. Cambridge, MA: MIT Press.

Amalric, M., & Dehaene, S. (2016). Origins of the brain networks for advanced mathematics in expert mathematicians. *Proceedings of the National Academy of Sciences of the United States of America.* 113:4909–4917. doi:10.1073/pnas.1603205113

Andres, M., Michaux, N., & Pesenti, M. (2012). Common substrate for mental arithmetic and finger representation in the parietal cortex. *NeuroImage*, 62(3), 1520–1528.

Ansari, D. (2008). Effects of development and enculturation on number representation in the brain. *Nature Reviews Neuroscience* 9, 278-291.

Ansari, D., & Dhital, B. (2006). Age-related changes in the activation of the intraparietal sulcus during nonsymbolic magnitude processing: an event-related functional magnetic resonance imaging study. *Journal of Cognitive Neuroscience. 18*, 1820–1828.

Ansari, D., Lyons, I. M., van Eimeren, L., & Xu, F. (2007). Linking visual attention and number processing in the brain: The role of the temporo-parietal junction in small and large symbolic and nonsymbolic number comparison. *Journal of Cognitive Neuroscience*, 19, 1845–1853.

Arp, S., Taranne, P., & Fagard, J. (2006). Global perception of small numerosities (subitizing) in cerebral-palsied children. *J Clin Exp Neuropsychol*, 28(3), 405–419.

Ashkenazi, S., Mark-Zigdon, N., & Henik, A. (2013). Do subitizing deficits in developmental dyscalculia involve pattern recognition weakness? *Developmental Science* 16:35–46.

Ashcraft, M. H. (1992). Cognitive arithmetic: A review of data and theory. *Cognition*, 44, 75–106.

Atkinson, J., Campbelle, F.W., & Francis, M.R. J. (1976). The magic number 4 plus or minus 0: A new look at visual numerosity judgements. *Perception*, S 327-334.

Aunger, R. (ed). (2001). *Darwinizing Culture: The status of memetics as a science*. Oxford: Oxford University Press.

Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in 5-month-old infants. *Cognition*, 20, 191–208.

Balakrishnan, J.D., &, Ashby, F.G. (1991). Is subitizing a unique numerical ability? *Perception and Psychophysics* 50: 555–64.

Baker, G.P. & Hacker, P.M.S.(1989). Frege's anti-psychologism. in M.A. Notturno (Ed.), *Perspectives on Psychologism*.(pp.74–127) Leiden: Brill.

Ball, B. (2016). On Representational Content and Format in Core Numerical Cognition. *Philosophical Psychology*, 1–21. doi:10.1080/09515089.2016.1263988.

Barner, D., Thalwitz, D., Wood, J. N., & Carey, S. (2007). Children's ability to distinguish "one" from "more than one" and the acquisition of singular-plural morpho-syntax. *Developmental Science*, 10(3), 365–373.

Barner, D., Wood, J. N., Hauser, M., & Carey, S. (2008). Evidence for a non- linguistic distinction between singular and plural sets in rhesus monkeys. *Cognition* 107 (2008) 603–622. doi:10.1016/j.cognition.2007.11.010

Barner, D., Chow, K., & Yang, S. (2009). Finding one's meaning: A test of the relation between quantifiers and integers in language development. *Cognitive*

Psychology, 58, 195-219.

Barner, D., Alvarez, G., Sullivan, J., Brooks, N., Srinivasan, M. and Frank, M. C. (2016), Learning Mathematics in a Visuospatial Format: A Randomized, Controlled Trial of Mental Abacus Instruction. *Child Development*, 87: 1146–1158. doi: 10.1111/cdev.12515

Barth, H. C., Kanwisher, N., & Spelke, E. (2003). The construction of large number representation in adults. *Cognition* 86, 201–221.

Basalla, G. (1988). *The evolution of technology*. New York: Cambridge University Press.

Beck, J. (2016). Can bootstrapping explain concept learning? *Cognition* 158, 110–121. http://dx.doi.org/10.1016/j.cognition.2016.10.017

Bell, J. L.(2018).Continuity and infinitesimals. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Summer 2017 Edition). Retrieved from https://plato.stanford.edu/archives/sum2017/entries/continuity/

Beller, S., & Bender, A. (2008). The limits of counting: Numerical cognition between evolution and culture. *Science*, 319, 213–215.

Benacerraf, P. (1965). What numbers could not be. *Philosophical Review* 74: 47–73.

Benacerraf, P. (1973). Mathematical truth. *Journal of Philosophy*, 70, 661–679.

Benacerraf, P., & Putnam, H. (1984). *Philosophy of Mathematics: Selected Readings. 2 edition.* York: Cambridge University Press.

Beran, M. J. (2004). Long-term retention of the differential values of Arabic numerals by chimpanzees (Pan troglodytes). *Animal Cognition* 7(2), 86–92.

Bernays, P. (1935). On Platonism in Mathematics. In P. Benacerraf & H. Putnam (Eds.) *Philosophy of Mathematics: Selected Readings. 2 edition.* York: Cambridge University Press.

Block, N.J. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy, 10*, 615–678.

Block, N. J. (1987). Functional role and truth conditions. *Proceedings of the Aristotelian Society*, 61, 157–181.

Boden, M. (2004). *The Creative Mind: Myths and Mechanisms* (2nd ed.). London:

Routledge.

Boysen, S. T., Bernston, G. G., Hannan, M. B., & Cacioppo, J. T. (1996). Quantity-based interfer- ence and symbolic representations in chimpanzees (Pan troglodytes). *J Exp Psychol Anim Behav Proces,* 22(1), 76–86.

Brannon, E. M. (2002). The development of ordinal numerical knowledge in infancy. *Cognition* 83, 223–240.

Brannon, E.M. (2005). The independence of language and mathematical reasoning. *Proceedings of the National Academy of Sciences of the United States of America* 102(9):3177–3178.

Brannon, E.M. & Terrace, H.S. (1998) Ordering of the numerosities 1 to 9 by monkeys. *Science* 282, 746–749.

Brown, R. (1973). *A first language: The early stages*. Cambridge, MA: Harvard University Press.

Brysbaert, M. (1995). Arabic Number Reading: On the Nature of the Numerical Scale and the Origin of Phonological Recoding. *Journal of Experimental Psychology: General* 124. 434-452. 10.1037/0096-3445.124.4.434.

Buckley, P.B., & Gillman, C.B. (1974). Comparison of digits and dot patterns. *Journal of Experimental Psychology*, 103, 1131-1136.

Bugden, S. & Ansari, D. (2011). Individual differences in children's mathematical competence are related to the intentional but not automatic processing of Arabic numerals. *Cognition* 118(1):32-44. doi: 10.1016/j.cognition.2010.09.005.

Bugden, S. & Ansari, D. (2016). Probing the nature of deficits in the 'approximate number system' in children with persistent developmental dyscalculia. *Developmental Science* 19(5):817–33. doi: 10.1111/desc.12324.

Burr, D.C., Turi, M., Anobile, G. (2010). Subitizing but not estimation of numerosity requires attentional resources. *Journal of Vision* 10(6):20, 1–10.

Butterworth, B. (1999). *The Mathematical Brain*. London: Macmillan.

Butterworth, B., & Reeve, R. (2008). Verbal counting and spatial strategies in numerical tasks: Evidence from indigenous Australia. *Philosophical Psychology* 21(4), 443–458.

Butterworth, B., Reeve, R., Reynolds, F., & Lloyd, D. (2008). Numerical thought with and without words: Evidence from indigenous Australian children. *Proceedings of the National Academy of Sciences of the USA* 105 (35), 13179-13184. https://doi.org/10.1073/pnas.0806045105

Campbell J.I.D., & Epp, L. J. (2004). An encoding-complex approach to numerical cognition in Chinese-English bilinguals. *Canadian Journal of Experimental Psychology*, 58, 229–244. doi: 10.1037/ h0087447 PMID: 15648727

Campbell, J.I.D. (2015). How abstract is arithmetic? In R. Cohen Kadosh & A. Dowker (Eds.), *The Oxford Handbook of Numerical Cognition*. Oxford: Oxford University Press.

Campbell J.I.D. & Clark J.M. (1988). An encoding-complex view of cognitive number processing: comment on McCloskey, Sokol, and Goodman (1986). *Journal of Experimental Psychology*: General, 117 (2), 204–14.

Cantlon, J.F., & Brannon, E. M. (2006). Shared system for ordering small and large numbers in monkeys and humans. *Psychological Science*, 17(5), 401–406.

Cantlon, J.F., Platt, M.L., and Brannon, E.M. (2009). Beyond the number domain. *Trends in Cognitive Science* 13(2): 83–91. doi: S1364-6613(08)00259-3

Capaldi, E. J., & Miller, D. J. (1988). Counting in rats: Its functional signi cance and the indepen- dent cognitive processes that constitute it. *Journal of Experimental Psychology: Animal Behavior Processes*, 14, 3–17.

Cappelletti, M., Gessaroli, E., Hithersay, R., Mitolo, M., Didino, D., Kanai, R. Cohen Kadosh, R. & Walsh, V. (2013). Transfer of Cognitive Training across Magnitude Dimensions Achieved with Concurrent Brain Stimulation of the Parietal Lobe. *Journal of Neuroscience* 11 33 (37) 14899-14907. DOI: 10.1523/JNEUROSCI.1692-13.2013

Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.

Carey, S. (2011). Précis of The Origin of Concepts. *Behavioral and brain sciences* 34 (3), 113-124.doi:10.1017/S0140525X10000919

Carter, J., & Palermos, S. (2016). Is Having Your Computer Compromised a Personal Assault? The Ethics of Extended Cognition. *Journal of the American Philosophical Association, 2*(4), 542-560. doi:10.1017/apa.2016.28

Chalmers, D.J. (2008). Foreword. In Clark, A. *Supersizing the mind.* (pp.ix-xxiv). Oxford: OUP.

Castelli, F., Glaser, D.E., and Butterworth, B. (2006). Discrete and analogue quantity processing in the parietal lobe: a functional MRI study. *Proceedings of the National Academy of Sciences of the USA 103*, 4693–4698.

Castronovo, J., & Crollen, V. (2011). Numerical comparison of two-digit numbers: How differences at encoding can involve differences in processing, *Journal of Cognitive Psychology* 23:1, 8-17, DOI: 10.1080/20445911.2011.445985

Castronovo, J., & Seron, X. (2007). Numerical estimation in blind subjects: Evidence of the impact of blindness and its following experience. *Journal of Experimental Psychology: Human Perception and Performance*, 33(5), 1089-1106. http://dx.doi.org/10.1037/0096-1523.33.5.1089

Cattell, J. M. (1886). Über die Trägheit der Netzhaut und des Sehcentrums. *Philosophische Studien*, 3, 94-127.

Cavanagh, P. (2011). Visual cognition. *Vision Research* 13, 1538-1551.

Charbonneau, M. (2015). All innovations are equal, but some more than others: (re)integrating modification processes to the origins of cumulative culture. *Biological Theory*, 10(4), 322–335.

Charbonneau, M. (2016). Modularity and recombination in technological evolution. *Philosophy & Technology*, 29: 373-392.

Cheries, E. W., Feigenson, L., Scholl, B. J., & Carey, S. (2005). Cues to object persistence in infancy: Tracking objects through occlusion and implosion. *Journal of Vision* 5(8):352. doi: 10.1167/5.8.352.

Chen, W.J., & Krajbich, I. (2017). Computational modeling of epiphany learning. *Proceedings of the National Academy of Sciences* May 2017, 114 (18) 4637-4642. DOI: 10.1073/pnas.1618161114.

Chesney, D.L., & Haladjian, H.H. (2011). Evidence for a shared mechanism used in multiple-object tracking and subitizing. *Attention Perception Psychophysics*, 73:2457–2480. DOI: 10.3758/s13414-011-0204-9

Cheung, P., Rubenson, M., & Barner, D. (2017). To infinity and beyond: Children generalize the successor function to all possible numbers years after learning to count. *Cognitive Psychology*, 92, 22-36.

Cheung, P., & Le Corre, M. (2018). Parallel individuation supports numerical comparisons in preschoolers. *Journal of Numerical Cognition*, 4 (2), 380-409.

Chochon, F., Cohen, L., van de Moortele, P.F., and Dehaene, S. (1999). Differential contributions of the left and right inferior parietal lobules to number processing. *Journal of Cognitive Neuroscience*. *11*, 617–630.

Chrisomalis, S. (2010). *Numerical notation: A comparative history*. Cambridge: Cambridge University Press.

Church, R. M., & Meck, W. H. (1984). The numerical attribute of stimuli. In H. L. Roitblat, T. G. Bever & H. S. Terrace (Eds.), *Animal Cognition* (pp. 445–464). Hillsdale N.J.: Erlbaum.

Cipolotti, L., Butterworth, B., & Denes, G. (1991). A specific deficit for numbers in a case of dense acalculia. *Brain* 114, 2619–2637.

Claidière, N., Scott-Phillips, T.C., and Sperber, D. (2014). How Darwinian is cultural evolution? *Philosophical Transactions of the Royal Society B: Biological Sciences* 369, no. 1642. doi:10.1098/rstb.2013.0368

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36 (03), 181-204. 10.1017/S0140525X12047

Clark, A. (1998). Twisted Tales: Causal Complexity and Cognitive Scientific Explanation. In F. Keil & R.A.Wilson. (Eds). *Explanation and Cognition* (p. 145-166). Cambridge, MA: MIT Press.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58: 7–19.

Clearfield, M. W., & Mix, K. S. (1999). Number versus contour length in infants' discrimination of small visual sets. Psychological Science, 10, 408–411.

Cohen L. B., & Marks K. S. (2002). How infants process addition and subtraction events. *Developmental Science* 5, 186–20110.1111/1467-7687.00220

Cohen Kadosh, R., Tzelgov, J., and Henik, A. (2008). A synesthetic walk on the number line: The size effect. *Cognition*, 106, 548-557

Cohen Kadosh, R. & Walsh, V. (2009). Numerical representation in the parietal lobes: Abstract or not abstract? *Behavioral and Brain Sciences*, 32(3–4), 313–28. doi: 10.1017/S0140525X09990938 PMID: 19712504

Cohen Kadosh, R. & Dowker, A. (Eds.) (2015) *The Oxford Handbook of Numerical*

Cognition. Oxford: Oxford University Press.

Coolidge, F. L., & Overmann, K.A. (2012). Numerosity, Abstraction, and the Emergence of Symbolic Thinking. *Current Anthropology* 53 (2). 204-25. doi:10.1086/664818

Corbetta, M., & Schulman, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience* 3, 201–215. doi:10.1038/nrn755

Cordes, S., Gelman, R., Gallistel, C. R., & Whalen, J. (2001). Variability signatures distinguish verbal from nonverbal counting for both large and small numbers. *Psychonomic Bulletin and Review,* 8(4), 698–707.

Cordes, S., & Gelman, R. (2005) The young numerical mind: When does it count? In Campbell, J. I. (Ed.). *Handbook of Mathematical Cognition*. New York: Psychology Press

Cordes, S., & Meck, W. H. (2013). Ordinal Judgments in the Rat: An Understanding of Longer and Shorter for Suprasecond, but Not Subsecond, Durations. *Journal of Experimental Psychology: General.* doi: 10.1037/a0032439

Crollen, V., Seron, X., & Noel, M.-P. (2011). Is finger-counting necessary for the development of arithmetic abilities? *Frontiers in Psychology*, 2(242), 1–3.

Cutini, S., Scatturin, P., Basso Moro, S., & Zorzi, M. (2013) Are the neural correlates of subitizing and estimation dissociable? An fNIRS investigation. *NeuroImage* 85 391–399. http://dx.doi.org/10.1016/j.neuroimage.2013.08.027

Crossley, J. N. (1987). *The emergence of number*. Singapore: World Scientific.

Darwin, C. (1871), *The descent of man and selection in relation to sex*. 2 vols. London: John Murray.

Dawkins, R. (1976/2006). *The Selfish Gene*. Oxford: OUP.

Dawkins, R. (1983). *The extended Phenotype*. Oxford: OUP.

Day, R.L., Laland, K.N., Odling-Smee, J. (2003). Rethinking Adaptation: the niche-construction perspective. *Perspectives in Biology and Medicine* 46(1):80–95.

De Cruz, H,. Hansjörg, N,. & Schlimm, D. (2010). The cognitive basis of arithmetic, in B. Löwe & T. Müller (Eds.) *Philosophy of Mathematics: Sociological Aspects and Mathematical Practice.*(pp. 59–106.). College Publications, London. Texts in Philosophy 11.

De Cruz, H. (2006) Why are some numerical concepts more successful than others? An evolutionary perspective on the history of number concepts. *Evolution and Human Behavior* 27:306–23.

De Cruz, H. (2007). *Innate ideas as a naturalistic source of mathematical knowledge: Towards a Darwinian approach to mathematics.* PhD Dissertation. Brussels: Vrije Universiteit Brussel.

De Cruz, H. (2008). An Extended Mind Perspective on Natural Number Representation. *Philosophical Psychology* 21, no. 4: 475–90.

De Cruz, H. (2016). Numerical cognition and mathematical realism. *Philosophers' Imprint*, 16, 1-13

De Smedt, B., Noël, M., Gilmore, C., & Ansari, D. (2013). How do symbolic and non-symbolic numerical magnitude processing skills relate to individual differences in children's mathematical skills? A review of evidence from brain and behavior. *Trends in Neuroscience and Education, 2, 48-55.*

Dehaene, S. (1992): Varieties of numerical abilities. *Cognition*, 44, 1-42.

Dehaene, S. (1996). The organization of brain activations in number comparison: Event-related potentials and the additive-factors methods. *Journal of Cognitive Neuroscience*, 8, 47–68.

Dehaene, S. (1997/2011) *The Number Sense: How the Mind Creates Mathematics.* New York: Oxford University Press.

Dehaene, S. (2003). The neural basis of the Weber-Fechner law : A logarithmic mental number line. *Trends in Cognitive Sciences*, 7, 145–147.

Dehaene, S. (2007). Symbols and quantities in parietal cortex : Elements of a mathematical theory of number representation and manipulation. In P. Haggard & Y. Rossetti (Eds.), *Attention & Performance XXII. Sensori-motor foundations of higher cognition.* (pp. 527–574). Cambridge, Mass : Harvard University Press.

Dehaene, S. (2009) *Reading in the Brain.* New York: Penguin.

Dehaene, S. (2014). Reading in the Brain Revised and Extended: Response to Comments. *Mind & Language*, Vol. 29, No. 3, pp. 320–335.

Dehaene, S., & Akhavein, R. (1995). Attention, automaticity and levels of representation in number processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21,314–326.

Dehaene, S., & Changeux, J.P. (1993). Development of elementary numerical abilities: A neuronal model. *Journal of Cognitive Neuroscience*, 5, 390–407.

Dehaene, S., & Changeux, J.P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron* 70, 200–227.

Dehaene, S., & Cohen, L. (1994). Dissociable mechanisms of subitizing and counting: Neuropsychological evidence from simaltanagnosic patients. *Journal of Experimental Psychology: Human Perception & Performance*, 20, 958-975.

Dehaene, S. & Cohen, L. (2007).Cultural recycling of Cortical maps. *Neuron*, 56 (2), 384-398. 10.1016/j.neuron.207.10.04

Dehaene, S., Dupoux, E., & Mehler, J. (1990). Is numerical comparison digital: Analogical and Symbolic effects in two-digit number comparison. *Journal of Experimental Psychology: Human Perception and Performance,* 16, 626–641.

Dehaene, S., Bossini, S., & Giraux, P. (1993). The mental representation of parity and number magnitude. *Journal of Experimental Psychology: General, 122*(3), 371-396.

Dehaene, S., Spelke, E., Pinel, P., Stanescu, R., and Tsivkin, S. (1999). Sources of mathematical thinking: Behavioral and brain-imaging evidence, *Science*, vol. 284, p.970- 974.

Demeyere, N., Lestou, V., & Humphreys, G. W. (2010). Neuropsychological evidence for a dissociation in counting and subitizing. *Neurocase*, 16, 1–19.

Demeyere N., Rotshtein, P. & and Humphreys, G.W. (2012). The Neuroanatomy of Visual Enumeration: Differentiating Necessary Neural Correlates for Subitizing versus Counting in a Neuropsychological Voxel-based Morphometry Study. *Journal of Cognitive Neuroscience* 24:4, 948-964

van Dijck, J.-P., Ginsburg, V., Girelli, L., & Gevers, W. (2015). Linking numbers to space: from the mental number line towards a hybrid account. In R. Cohen Kadosh & A. Dowker (Eds.), *The Oxford handbook of numerical cognition* (pp. 89–105). Oxford, UK: Oxford University Press.

Donlan, C., Cowan, R., Newton, E., and Lloyd, D. (2007). The role of language in mathematical development: Evidence from children with specific language impairments. *Cognition.* 103(1):23–33.

Drummond, J.J. (1985). Frege and Husserl: Another look at the influence. *Husserl Studies* 2:245-265.

Dummett, M. (2001). Preface. In E. Husserl, *Logical Investigations*, J.N Findlay (trans). Vol 1.

Duncan, J., & Humphreys, G.W. (1989) Visual search and stimulus similarity. *Psychological Review* 96 (3) 433–58.

Dutilh Novaes, C. (2013). Mathematical reasoning and external symbolic systems. *Logique & Analyse* 56 (21): 45-65.

Eger, E., Sterzer, P., Russ, M.O., Giraud, A.L., and Kleinschmidt, A. (2003). A supramodal number representation in human intraparietal cortex. *Neuron 37*, 719–725.

Egeth, H.E., Leonard, C.J., & Palomares, M. (2008) The role of attention in subitizing: Is the magical number 1? *Visual Cognition* 16: 463–473.

Emlen, S. (1975). The stellar orientation system of a migratory bird. *Scientific American*, 233: 102–111.

Everett, C. (2017). *Numbers and the making of us*. Cambridge, Mass.: Harvard University Press.

Everett, D. L. (2005). Cultural constraints on grammar and cognition in Pirahã. *Current Anthropology* 46, 621–646.

Fabry, R. E. (2016). Enriching the notion of enculturation: Cognitive integration, predictive processing, and the case of reading acquisition: A commentary on Richard Menary. In T. Metzinger & J. M. Windt (Eds.), *Open MIND: Philosophy and the mind sciences in the 21 century* (Vol 2, pp. 1017–1039). Cambridge, Mass.: MIT Press.

Fayol, M., & Seron,X.(2005). About numerical representations: insights from neuropsychological, experimental and developmental studies. p. 3–22. In Campbell, J. I. (Ed.). Handbook of Mathematical Cognition. New York: Psychology Press.

Feigenson, L., Carey, S., & Spelke, E. (2002). Infants' discrimination of number vs. continuous extent. Cognitive Psychology, 44, 33–66

Feigenson, L., Dehaene, S., and Spelke, E. (2004) Core systems of number. *Trends in Cognitive Sciences* 8:307–14.

Feigenson, L. (2007). The equality of quantity. *Trends in Cognitive Sciences*, 11(5), 185-187.

Ferreirós, J. (2015). *Mathematical knowledge and the interplay of practices*. Princeton: Princeton University Press.

Flombaum, J.L., Kundey, S.M., Santos, L.R., & Scholl, B.J. (2004). Dynamic object individuation in rhesus macaques: a study of the tunnel effect. *Psychological Science*. 15(12):795-800.

Flombaum, J. I., Junge, J., & Hauser, M. D. (2005). Rhesus monkeys (Macaca mulatta) spontaneously compute large number addition operations. *Cognition*, 97(3), 315–325.

Fodor, J.A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. New York: Oxford University Press.

Fodor, J.A. (2010). Woof, woof. *Times Literary Supplement*, October 8: 7–8.

Fornaciai, M., Cicchini, G. M., & Burr, D. C. (2016). Adaptation to number operates on perceived rather than physical numerosity. *Cognition* 151, 63–67. doi: 10.1016/j.cognition.2016.03.006

Fornaciai, M., & Park, J. (2017) Distinct Neural Signatures for Very Small and Very Large Numerosities. *Front. Hum. Neurosci.* 11:21. doi: 10.3389/fnhum.2017.00021

Foster, G. M. (1944). *A summary of Yuki culture*. Berkeley, CA: University of California Press.

Frege, G. (1884/1980). *The foundations of Arithmetic : A logico-mathematical enquiry into the concept of number*. Evanston, Ill: Northwestern University Press.

Frege, G. (1893/1964), Grundgesetze der Arithmetik. Translated (in part) by M. Furth as The Basic Laws of Arithmetic, Berkeley, CA: University of California Press.

Frank, M., Everett, D., Fedorenko, E., & Gibson, E. (2008). Number as a cognitive technology: Evidence from Pirahã language and cognition. *Cognition* 108: 819-824.

Gallistel, C. R. (1990). *The Organization of Learning*. Cambridge (Mass.): MIT Press.

Gallistel, C. R. (2007). Commentary on Le Corre & Carey. *Cognition* 105, 439–445.

Gallistel, C.R. & Gelman, R. (1991) Subitizing: the preverbal counting process. In *Memories, Thoughts and Emotions: Essays in Honor of George Mandler* (Kessen, W. et al., Eds), pp. 65–81, Erlbaum

Gallistel, C. R., & Gelman, R. (1992). Preverbal and verbal counting and computation. *Cognition*, 44, 43–74.

Gallistel, C. R., & Gelman, R. (2000). Non-verbal numerical cognition: from reals to integers. *Trends in Cognitive Sciences* 4(2):59-65.

Gallistel, C. R., Gelman, R., & Cordes, S. (2006) The cultural and evolutionary history of the real numbers. In: *Evolution and culture*, in . S. C. Levinson and P. Jaisson (Eds.). (pp. 247–74). Cambridge, MA: MIT Press.

Galton, F. (1880). Visualised numerals. *Nature* 21:252–256.

Gebuis, T. & Reynvoet, B. (2012a). Continuous visual properties explain neural responses to nonsymbolic number. *Psychophysiology* 49(11):1481–91. doi: 10.1111/j.1469-8986.2012.01461.

Gebuis, T. & Reynvoet, B. (2012b). The interplay between nonsymbolic number and its continuous visual properties. *Journal of Experimental Psychology: General* 141(4):642–48. doi: 10.1037/a0026218.

Gebuis, T., Gevers, W. & Cohen Kadosh, R. (2014). Topographic representation of high-level cognition: Numerosity or sensory processing? *Trends in Cognitive Sciences* 18(1):1–3. doi: 10.1016/j.tics.2013.10.002.

Gebuis, T., Cohen Kadosh, R. & Gevers, W. (2016) Sensory-integration system rather than approximate number system underlies numerosity processing: A critical review. *Acta Psychologica* 171:17–35.

Gelman, R. and Butterworth, B. (2005). Number and language: How are they related? *Trends in Cognitive Sciences* 9, 6–10.

Gelman, R., & Gallistel, C. R. (1978). *The child's understanding of numb*er. Cambridge: Harvard University Press/MIT Press.

Gilmore, C.K., McCarthy, S. E., and Spelke, E.S. (2007). Symbolic arithmetic knowledge without instruction. *Nature 447*, 589–591.

Gliksman, Y., Weinbach, N., & Henik, A. (2016). Alerting cues enhance the subitizing process. *Acta Psycholica* 170 139–145. doi: 10.1016/j.actpsy.2016.06.013

Göbel S. M., Shaki S., Fischer M. H. (2011). The cultural number line: a review of cultural and linguistic influences on the development of number processing. *Journal of Cross-Cultural Psychology* 42, 543–56510.1177/0022022111406251

Godden, D. (2005) Psychologism in the Logic of John Stuart Mill: Mill on the Subject Matter and Foundations of Ratiocinative Logic. *History and Philosophy of Logic* 26 (2):115-143.

Gordon, P. (2004). Numerical cognition without words: Evidence from Amazonia. *Science* 306:496–99.

Green, C. S., Bavelier, D. (2004). Does action video game play really enhance the number of items that can be simultaneously attended? [Abstract]. *Journal of Vision*,4(8): 632, 632a, http://journalofvision.org/4/8/632/, doi:10.1167/4.8.632.

Greer, B. (2004). The growth of mathematics through conceptual restructuring. *Learning and Instruction*, 14, 541–548.

Guha, A., Jereen, A., Degutis, J. & Wilmer, J. (2014). No action video game training effects for multiple object tracking or mental rotation. *Journal of Vision*. 14. 1323. 10.1167/14.10.1323.

Hadamard, J. (1945). *The Mathematician's Mind: The Psychology of Invention in the Mathematical Field*. Princeton : Princeton University Press.

Haith M. M. (1998). Who put the cog in infant cognition? Is rich interpretation too costly. *Infant Behavior and Development*. 21, 167–17910.1016/S0163-6383(98)90001-7.

Haladjian, H.H., & Pylyshyn, Z. (2010). Enumerating by pointing to locations: a new method for measuring the numerosity of visual object representations. *Attention Perception & Psychophysics* 73:303–308 DOI 10.3758/s13414-010-0030-5.

Haladjian, H. H., Singh, M., Pylyshyn, Z. W., & Gallistel, C. R. (2010). The encoding of spatial information during small-set enumeration. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd annual conference of the cognitive science society* (pp. 2839–2844). Austin, TX: Cognitive Science Society.

Halberda, J., & Feigenson, L. (2008). Developmental change in the acuity of the" Number Sense": The Approximate Number System in 3-, 4-, 5-, and 6-year-olds and adults. *Developmental Psychology* 44(5), 1457–1465.

Halberda, J., Ly, R., Wilmer, J. B., Naiman, D. Q., & Germine, L. (2012). Number sense across the lifespan as revealed by a massive Internet-based sample. *Proceedings of the National Academy of Sciences* 109(28), 11116–11120.

Haugeland, J. (1993). Mind embodied and embedded. In Yu-Houng H. Houng & J. Ho (eds.), *Mind and Cognition: 1993 International Symposium*. Academica Sinica. pp. 233-267

Hauser, M. D., & Carey, S. (1998). Building a cognitive creature from a set of primitives: Evolutionary and developmental insights. In D. Cummins & C. Allen (Eds.), The evolution of mind. (pp. 51–106). Oxford, UK: Oxford University Press.

Hauser, M. D., Carey, S., Hauser, L.B. (2000). Spontaneous number representations in semi-free-ranging rhesus monkeys. *Proceedings of the Royal Society of London B,* vol. 267, p. 829-833.

Hauser, M. D., & Carey, S. (2003). Spontaneous representations of small numbers of objects by rhesus macaques: examinations of content and format. *Cognitive Psychology*, 47(4), 367–401.

Heidelberger, M. (2004). *Nature From Within: Gustav Theodor Fechner and His Psychophysical Worldview.* Pittsburgh: University of Pittsburgh Press.

Henrich, J., Boyd, R. & Richerson, P.J. (2008).Five misunderstandings about cultural evolution. *Human Nature* 19: 119. https://doi-org.proxy.bibliotheques.uqam.ca:2443/10.1007/s12110-008-9037-1

de Hevia, M.D., Castaldi, E., Streri, A., Eger, E., & Izard, V. (2017). Perceiving numerosity from birth. *Behavioral and Brain Sciences*, 40, E169. doi:10.1017/S0140525X16002090

Hinrichs, J. V., Yurko, D. S., & Hu, J. M. (1981). Two-digit number comparison: use of place information. *Journal of Experimental Psychology: Human Perception And Performance* 7, 890–901.

Hubbard, E.M., Piazza, M., Pinel, P., Dehaene, S. (2005) Interactions between number and space in parietal cortex. *Nature Reviews Neuroscience* 6(6): 435–448

Hubbard, E., Piazza, M., Pinel, P., & Dehaene, S. (2009). Numerical and Spatial Intuitions: A Role for Posterior Parietal Cortex? In L. Tommasi, L. Nadel & M. A. Peterson (Eds.), *Cognitive Biology: Evolutionary and Developmental Perspectives on Mind, Brain and Behavior* (pp. 221–246). Cambridge, Mass.: MIT Press.

Hurford, J. R. (1987). *Language and Number*. Oxford: Basil Blackwell.

Husserl, E. (1900/2001). *Logical Investigations.* Translated by Findlay, J.N. London: Routledge

Hutchins, E. (1995). *Cognition in the wild.* Cambridge, Mass.: MIT Press.

Hyde, D.C., & Spelke, E.S. (2009) All numbers are not equal: An electrophysiological investigation of small and large number representations. *Journal of Cognitive Neuroscience*. 21(6):1039–1053.

Hyde D.C., & Spelke, E.S. (2011). Neural signatures of number processing in human infants: Evidence for two core systems underlying non-verbal numerical cognition. *Developmental Science* 14(2):360-71. doi: 10.1111/j.1467-7687.2010.00987.x.

Hyde, D. C., Khanum, S. & Spelke, E. S. (2014) Brief non-symbolic, approximate number practice enhances subsequent exact symbolic arithmetic in children. *Cognition* 131:92–107.

Hyde, D.C. & Mou, Y. (2017). Magnitude rather than number: More evidence needed. *Behavioral and Brain Sciences*, 40, e173

Ifrah, G. (1998). *The Universal History of Numbers*. London: The Harvil Press.

Inglis, M., Batchelor, S,. Gilmore, C., Watson, D.G. (2017). Is the ANS linked to mathematics performance? *Behavioral and Brain Sciences* 40:e174. doi: 10.1017/S0140525X16002120.

Ischebeck, A. (2003). Differences between digit naming and number word reading in a flanker task. *Memory and Cognition* 31 (4), 529-537.

Ischebeck, A. (2003) Differences between digit naming and number word reading in a flanker task. Memory and Cognition 31:529–37.

Ischebeck, A., Zamarian, L., Schocke, M., & Delazer, M. (2009). Flexible transfer of knowl- edge in mental arithmetic—an fMRI study. *NeuroImage*, 44, 1103–1112.

Ito, Y. & Hatta, T. (2003) Semantic processing of Arabic, Kanji, and Kana numbers: Evidence from interference in physical and numerical size judgments. *Memory and Cognition* 31(3):360–68.

Ito, Y., & Hatta, T. (2004). Spatial structure of quantitative representation of numbers: evidence from the SNARC effect. *Memory and Cognition* 32(4), 662–673.

Izard, V., and Dehaene, S. (2007). Calibrating the mental number line. *Cognition*, 106(3), 1221–1247. doi:10.1016/j.cognition.2007.06.004

Izard, V., Pica, P., Spelke, E.S., Dehaene, S. (2008) Exact Equality and Successor Function: Two Key Concepts on the Path towards Understanding Exact Numbers, *Philosophical Psychology*,21:4,491 — 505. doi: 10.1080/09515080802285354

Izard, V., Sann, C., Spelke, E. S. & Steri, A. (2009) Newborn infants perceive abstract numbers. *Proceedings of the National Academy of Sciences of the United States of America* 106(25):10382–85.

Jacob, S.N., & Nieder, A. (2009). Notation-independent representation of fractions in the human parietal cortex. *Journal of Neuroscience* 29, 4652–4657. doi:10.1523/jneurosci.0651-09.2009

Jennings, C. & Nanay, B. (2014). Action without Attention. *Analysis*. 76.10.1093/analys/anu096.

Jevons, W.S. (1871). The power of numerical discrimination. *Nature, 3*, 281-282.

Johnson, M. H., Bolhuis, J. J., & Horn, G. (1985). Interaction between acquired preferences and developing predispositions during imprinting. *Animal Behaviour*, 33, 1000–1006.

Johnson, S.P. (2010). How Infants Learn About the Visual World. *Cognitive Science* (34) 7. 1158-1184. doi:10.1111/j.1551-6709.2010.01127.x

Jordan, K.E., & Brannon, E. M. (2006).The multisensory representation of number in infancy. . *Proceedings of the National Academy of Sciences of the United States of America* 103(9): 3486–3489. doi: 10.1073/pnas.0508107103

Jordan, K,. Maclean, E., & Brannon, E. (2008). Monkeys match and tally quantities across senses. *Cognition* 108. 617-25. 10.1016/j.cognition.2008.05.006.

Joseph, J. S., Chun, M. M., & Nakayama, K. (1997). Attentional requirements in a "preattentive" feature search task. *Nature*, 387(6635), 805-807. http://dx.doi.org/10.1038/42940

Kaan, E. (2004). Direction effects in number word comparison: An event-related potential study. *Neuroreport* 16(16):1853-6.

Kahneman, D., Treisman, A. & Gibbs, B. J. (1992) The reviewing of object files: Object-specific integration of information. *Cognitive Psychology* 24:175– 219.

Káldy, Z., & Leslie, A. M. (2005). A memory span of one? Object identification in 6.5-month-old infants. *Cognition*, 97, 153–177.

Kaufman, E. L., Lord, M. W., Reese, T. W., & Volkmann, J. (1949). The discrimination of visual number. *American Journal of Psychology*, 62, 498-525.

Kaufmann L., & Nuerk H.C. (2008). Basic number processing deficits in ADHD: A broad examination of elementary and complex number processing skills in 9- to 12-year-old children with ADHD-C. *Developmental Science* 11:692–699.

Kellman, P. J., & Spelke, E. S. (1983). Perception of partly occluded objects in infancy. *Cognitive Psychology*, 15(4), 483–524.

Kitterle, F.L., and Selig, L.M. (1991). Visual field effects in the discrimination of sine-wave gratings. *Percept. Psychophys. 50*, 15–18.

Kirkpatrick, L. A. (2009). Between evolution and culture: Psychology at the nexus. In M.Schaller, S. J. Heine, A. Norenzayan, T. Yamagishi, and T. Kameda (Eds.), *Evolution, culture, and the human mind* (pp. 71-79) Mahwah, NJ: Lawrence Erlbaum.

Kirsh, D., & P. Maglio. (1992). Reaction and reflection in Tetris. In J. Hendler. (Ed.) *Artificial intelligence planning systems: Proceedings of the first annual conference AIPS,* San Mateo, CA: Morgan Kaufmann.

Knops, A., Viarouge, A., and Dehaene, S. (2009). Dynamic representations underlying symbolic and nonsymbolic calculation: evidence from the operational momentum effect. *Attention Perception Psychophysics* 71, 803–821. doi:10.3758/APP.71.4.803

Knops, A., Zitzmann, S., McCrink, K. (2013) Examining the presence and determinants of operational momentum in childhood. *Frontiers of Psychology* 4: 325

van Kerkhove, B., Van Bendegem, J. P., & De Vuyst, J. (Eds.) (2010). *Philosophical Perspectives on Mathematical Practice*. College Publications.

van Kerkhove, B & Van Bendegem, J.P (Eds). (2006) FULL ISSUE. *Foundations of Science* 11, no. 1-2.

Kitcher, P. (1984). *The nature of mathematical knowledge*. Oxford: Oxford University Press.

Klahr, D. (1973). Quantification processes. In W. G. Chase (Ed.), *Visual information* *processing* (pp. 3–34). New York: Academic Press.

Kline, M. (1980). *Mathematics: The Loss of Certainty*. New York: Oxford University Press.

Koehler, O. (1951). The ability of birds to count. *Bulletin of animal behavior*, 9, 41-45.

Kusch, M. (1995). *Psychologism: The Sociology of Philosophical Knowledge*. London ; New York: Routledge.

Lakatos, I. (1976). *Proofs and refutations: The logic of mathematical discovery.*Cambridge: Cambridge University Press.

Lakoff, G., and Núñez, R. E. (2000). *Where mathematics comes from: How the embodied mind brings mathematics into being*. New York: Basic Books.

Laland, K.N., Boogert, N., Evans, C. (2014) Niche construction, innovation and complexity. *Environmental Innovation and Societal Transitions* 11: 71-86. http://dx.doi.org/10.1016/j.eist.2013.08.003

Landy, D. & Goldstone, R.L. (2007). Formal notations are diagrams: evidence from a production task. *Memory and Cognition* 35(8), 2033–2040.

Lavie, N., & Tsal, Y. (1994). Perceptual Load as a Major Determinant of the Locus of Selection in Visual Attention. *Perception & Psychophysics* 56 (2): 183–97. doi:10.3758/BF03213897.

Lavie, N. Hirst, A. de Fockert J. W., & Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology General* 133(3):339-54.

Le Corre, M., & Carey, S. (2007). One, two, three, four, nothing more: How numerals are mapped onto core knowledge of number in the acquisition of the counting principles. *Cognition*, 105, 395–438.

Le Corre, M., & Carey, S. (2008). Why the verbal counting principles are constructed out of representations of small sets of individuals: A reply to Gallistel. *Cognition* 107 (2008) 650–662.

Leibovich, T. & Henik, A. (2013) Magnitude processing in non-symbolic stimuli. *Frontiers in Psychology* 4:375. doi: 10.3389/fpsyg.2013.00375.

Leibovich, T., & Ansari, D. (2016). The symbol-grounding problem in numerical cognition: A review of theory, evidence, and outstanding questions. *Canadian Journal of Experimental Psychology/Revue Canadienne De Psychologie Expérimentale*, 70, 12–23. doi:10.1037/cep0000070

Leibovich, T., Katzin, N., Harel, M., and Henik, A. (2017). From 'sense of number' to 'sense of magnitude' - The role of continuous magnitudes in numerical cognition. *Behavioral and Brain Sciences*. 1-62. doi:10.1017/S0140525X16000960

Lerner, Y., Hendler, T., Ben-Bashat, D., Harel, M., and Malach, R. (2001). A hierarchical axis of object processing stages in the human visual cortex. *Cerebral Cortex 11*, 287–297.

Leslie, A. M., Xu, F., Tremoulet, P. D., & Scholl, B. J. (1998). Indexing and the object concept: Developing "what" and "where" systems. Trends in Cognitive Science, 2, 10–18.

Lewontin, R.C. (1983). Gene, organism and environment. In: Bendall, D.S. (Ed.), *Evolution from Molecules to Men*. Cambridge: Cambridge University Press, pp. 273–285.

Libertus, M. E., Woldorff, M. G. & Brannon, E. M. (2007). Electrophysiological evidence for notation independence in numerical processing. *Behavioral Brain Function* 3(1).

Li, P., Le Corre, M., Shui, R., Jia, G., & Carey, S. (2003). Effects of plural syntax on number word learning: A cross-linguistic study. Paper presented at the 28thAnnual Boston University Conference on Language Development, Boston, MA.

Li, P., Ogura, T., Barner, D., Yang, S.-J., & Carey, S. (2009). Does the conceptual distinction between singular and plural sets depend on language? *Developmental Psychology*, 45(6), 1644-1653.http://dx.doi.org/10.1037/a0015553

Lipton, J. S., & Spelke, E. S. (2003). Origins of number sense: Large number discrimination in human infants. *Psychological Science*, 15(5), 396–401.

Lipton, J. S., & Spelke, E. (2004). Discrimination of large and small numerosities by human infants. *Infancy* 5(3), 271–290.

Lipton, J., & Spelke, E. (2005). Preschool Childrens Mapping of Number Words to Nonsymbolic Numerosities. *Child Development* 76 (5): 978-988.

Longo, M.R., & Lourenco, S.F. (2007). Spatial attention and the mental number line: Evidence for characteristic biases and compression. *Neuropsychologia* 45, 1400–

1407.

Lorenz, K. (1937). The companion in the bird's world. *Auk*, 54, 245–273.

Lourenco, S. F. (2015) On the relation between numerical and non-numerical magnitudes: Evidence for a general magnitude system. In: *Evolutionary origins and early development of number processing*. D. C. Geary, D. B. Berch & K. M. Koepke (Eds). (pp. 145–74.) London: Elsevier Academic Press.

Lukas, S. Krinzinger, H., Koch, I. & Willmes, K. (2014) Number representation: A question of look? The distance effect in comparison of English and Turkish number words, *The Quarterly Journal of Experimental Psychology,* 67:2, 260-270, DOI: 10.1080/17470218.2013.802002

Lyons, I.M., Ansari, D., and Beilock S.L. (2012). Symbolic Estrangement: Evidence Against a Strong Association Between Numerical Symbols and the Quantities They Represent. Journal of Experimental Psychology: General, 141 (4):635-641. doi: 10.1037/a027248

Macbeth, D. (2013). Writing Reason. *Logique & Analyse* 221, 25–44.

Macizo, P. & Herrera, A. (2010). Two-Digit Number Comparison: Decade-Unit and Unit-Decade Produce the Same Compatibility Effect With Number Words. *Canadian Journal of Experimental Psychology* 64, No. 1, 17–24.

Macizo, P. & Herrera, A. (2014). Facilitation and interference in the color-naming task. *Memory and Cognition* 42 (6) : 843–853 DOI 10.3758/s13421-014-0413-2

Mack, A., & Rock, I. (1998). *Inattentional Blindness*, Cambridge, MA: MIT Press.

MacLeod, C. M., & Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*(1), 126-135.

Maddy, P.(1996). Set Theoretic Naturalism. *Journal of Symbolic Logic* 61(2): 490–514.

Maddy, P. (1997). *Naturalism in Mathematics*. Oxford: Clarendon.

Maddy, P. (1998). How to be a Naturalist about Mathematics. In H.G. Dales & G Oliveri (Eds.), *Truth in Mathematics*. (pp.161–80). Oxford: Clarendon.

Malafouris, L. (2010). Grasping the concept of number: how did the sapient mind move beyond approximation? In *The archaeology of measurement: comprehending* *heaven, earth and time in ancient societies.* C. Renfrew and I. Morley, (Eds). pp. 35–42. Cambridge: Cambridge University Press.

Mandler, G., & Shebo, B. J. (1982). Subitizing: An analysis of its component processes. *Journal of Experimental Psychology: General*, 111, 1–22. doi:10.1037/0096-3445.111.1.1

Maloney, E. A., Risko, E. F., Ansari, D., & Fugelsang, J. (2010). Mathematics anxiety affects counting but not subitizing during visual enumeration. *Cognition*, 114(2), 293–297.

Marr, D. (1982). *Vision.* San Francisco: W. H. Freeman and Company.

Matsuzawa, T. (1985). Use of numbers by a chimpanzee. *Nature*, 315, 57–59.

Matsuzawa, T. (2009).Symbolic representation of number in chimpanzees. *Current opinion in Neurobiology* 19:92-98.

Mazza, V.(2017). Simultanagnosia and object individuation, *Cognitive Neuropsychology,* DOI: 10.1080/02643294.2017.1331212

McCloskey, M. (1992): Cognitive mechanisms in numerical processing: evidence from acquired dyscalculia. *Cognition*, 44, 107-157.

McCloskey M., Caramazza A., & Basili A. (1985). Cognitive mechanisms in number processing and calculation: Evidence from dyscalculia. Brain and Cognition, 4, 171–96. PMID: 2409994

McCloskey, M. & Macaruso, P. (1995) Representing and using numerical information. *American Psychologist* 50:351–63.

McCrink, K., & Wynn, K. (2004). Large-number addition and subtraction by 9-month-old infants. *Psychological Science*, 15, 776–781. doi:10.1111/j.0956-7976.2004.00755

Mechner, F. (1958). Probability relations within response sequences under ratio reinforcement. *Journal of the Experimental Analysis of Behavior*, 1, 109–121.

Mechner, F., & Guevrekian, L. (1962). Effects of deprivation upon counting and timing in rats. *Journal of the Experimental Analysis of Behavior*, 5, 463–466.

Meck, W. H., & Church, R. M. (1983). A mode control model of counting and timing processes. *Journal of Experimental Psychology: Animal Behavior Processes*, 9, 320–334.

Menary, R. (2007). *Cognitive Integration: Mind and Cognition Unbounded.* Basingstoke: Palgrave Macmillan.

Menary, R. (ed). (2010a). *The Extended Mind.* Cambridge: MIT Press.

Menary, R. (2010b). Dimensions of mind. *Phenomenology and the Cognitive Sciences*, 9, 561-578. 10.107/s1097-010-9186-7

Menary, R. (2015a). Mathematical Cognition - A Case of Enculturation. In T. Metzinger & J. M. Windt (Eds.) *Open MIND*. Frankfurt a. M., GER: MIND Group.

Menary, R. (2015b). What? Now. Predictive Coding and Enculturation – A Reply to Regina E. Fabry. In T. Metzinger & J. M. Windt (Eds.) *Open MIND*. Frankfurt a. M., GER: MIND Group. doi:10.15502/9783958571198

Menninger, K. (1969). *Number words and number symbols: A cultural history of numbers*. Cambridge, MA: MIT Press.

Mill, J.S., 1865. An Examination of Sir William Hamilton's Philosophy, in J.M. Robson (Ed.), *The Collected Works of John Stuart Mill, Volume 9*, London and Toronto: Routledge and University of Toronto Press, 1979.

Mix, K. S., Levine, S. C., & Huttenlocher, J. (1997). Numerical abstraction in infants: Another look. *Developmental Psychology*, 33, 423–428.

Mix, K. S., Huttenlocher, J., & Levine, S. C. (2002). Multiple cues for quantification in infancy: Is number one of them? *Psychological Bulletin*, 128(2), 278-294. DOI: 10.1037//0033-2909.128.2.278

Moeller, K., Pixner,S., Kaufmann,L. & Nuerk, H.C. (2009) Children's early mental number line: logarithmic or decomposed linear? *Journal of Experimental Child Psychology*, 103, 503-515

Moeller K., Fischer U., Link T., Wasner M., Huber S., Cress U., et al. . (2012). Learning and development of embodied numerosity. *Cogn. Process*. 13, S271–S274. 10.1007/s10339-012-0457-9.

Mohanty, J.J. (1982). *Husserl and Frege.* Bloomington: Indiana University Press.

Möhring, W., Frick, A., & Newcombe N.S. (2018). Spatial scaling, proportional thinking, and numerical understanding in 5- to 7-year-old children. *Cognitive Development* 45 (2018) 57–67. https://doi.org/10.1016/j.cogdev.2017.12.001

Moyer, R.S. & Landauer, T.K. (1967). The time required for judgments of numerical inequality. *Nature* 215: 1519-1520.

Myers, T. & Szűcs, D. (2015). Arithmetic Memory Is Modality Specific. *PloS one*. 10. e0145614. 10.1371/journal.pone.0145614.

Nieder, A., & Miller, E. K. (2003). Coding of cognitive magnitude: compressed scaling of numerical information in the primate prefrontal cortex. *Neuron 37(1):149-57.*

Nieder, A., & Miller, E. K. (2004). A parieto-frontal network for visual numerical information in the monkey. Proc. Natl. Acad. Sci. USA *101*, 7457–7462.

Nieder, A. (2005). Counting on neurons: the neurobiology of numerical competence. *Nature Reviews Neuroscience* 6(3), 177–190.

Nieder, A., & Merten, K. (2007). A labeled-line code for small and large numerosities in the monkey prefrontal cortex. *J. Neurosci.*, *27*(22), 5986–5993.

Nieder, A. & Dehaene, S. (2009).Representations of number in the brain. *Annual Review of Neuroscience*, 32, 185-208.

Nuerk, H.-C., Moeller, K., Klein, E., Willmes, K., and Fischer, M. H. (2011). Extending the Mental Number Line—A Review of Multi-Digit Number Processing. *Journal of Psychology* 219: 3–22.

Nuerk, H.-C., Moeller, K., & Willmes, K. (2015) Multi-digit number processing – overview, conceptual clarifications, and language influences. In R. Cohen Kadosh & A. Dowker (Eds.), *The Oxford Handbook of Numerical Cognition.* Oxford: Oxford University Press.

Núñez, R., Cooperrider, K., & Wassmann, J. (2012). Number Concepts without Number Lines in an Indigenous Group of Papua New Guinea. *PLoS ONE* 7(4): e35662. doi:10.1371/journal.pone.0035662

Núñez, R. E. (2011). No Innate Number Line in the Human Brain. *Journal of Cross-Cultural Psychology* 42: 651 doi: 10.1177/0022022111406097

Núñez, R. E. (2017). Is there really an evolved capacity for number? *Trends in Cognitive Sciences,* 21(6), 409-424.

Nys, J. & Content, A. (2012) Judgement of discrete and continuous quantity in adults: Number counts! *Quarterly Journal of Experimental Psychology* 65 (4):675–90. doi: 10.1080/17470218.2011.619661.

van Oeffelen, M. P., & Vos, P. G. (1982). A probabilistic model for the discrimination of visual number. *Perception & Psychophysics*, 32, 163–170.

Oksama, L., & Hyona, J. (2004). Is multiple object tracking carried out automatically by an early vision mechanism independent of higher-order cognition? An individual difference approach. *Visual Cognition*, 11(5), 631–71.

Olivers, C.N.L., Watson, D.G. (2008) Subitizing requires attention. *Visual Cognition* 16: 439–463.

Overmann, K. A., Wynn, T., & Coolidge, F.L. (2011). The prehistory of number concept. *Behavioral and Brain Sciences* 34 :3. 142-144.

Overmann, K. A. (2014). Finger-counting in the Upper Palaeolithic. *Rock Art Research* 31(01):63–80.

Overmann, K.A. (2015). Numerosity Structures the Expression of Quantity in Lexical Numbers and Grammatical Number. *Current Anthropology* 56 (5), 638-653.

Overmann, K. A. (2016). The role of materiality in numerical cognition. *Quaternary International*, 405, 42–51.

Oyama, S. (1985). *The ontogeny of information: Developmental systems and evolution*. New York: CUP.

Pagano, S., Lombardi, L., Mazza, V. (2014). Brain dynamics of attention and working memory engagement in subitizing. *Brain Research* 1543, 244–252.

Park, J. & Brannon, E. M. (2014). Improving arithmetic performance with number sense training: An investigation of underlying mechanism. *Cognition* 133:188–200.

Pepperberg, I. M., & Gordon, J. D. (2005). Number comprehension by a grey parrot (Psittacus erithacus), including a zero-like concept. *Journal of Comparative Psychology*, 119, 197–209.

Pepperberg, I. M. (2006). Grey parrot numerical competence: A review. *Animal Cognition*, 9, 377–391.

Priftis, K,. Zorzi, M., Meneghello, F., Marenzi, R., & Umiltà, C. (2006). Explicit versus implicit processing of representational space in neglect: Dissociations in accessing the mental number line. *Journal of Cognitive Neuroscience* 18(4): 680–688.

Philipse, H. (1989) Psychologism and the prescriptive function of logic. in M.A. Notturno (ed.), *Perspectives on Psychologism.* (pp. 58–74). Leiden: Brill.

Piaget, J. (1980). Language within cognition: Schemes of action and language learning. In M. Piatelli-Palmerini (Ed.), *Language and learning.* (pp. 163–183). Cambridge, MA: Harvard University Press.

Piazza, M., Mechelli, A., Butterworth, B., & Price, C. J. (2002). Are subitizing and counting implemented as separate or functionally overlapping processes? *Neuroimage*, 15(2), 435–446.

Piazza, M., Giacomini, E., Le Bihan, D., & Dehaene, S. (2003). Single-trial classification of parallel pre-attentive and serial attentive processes using functional magnetic resonance imaging. Proc R Soc Lond B Biol Sci, 270(1521), 1237–1245.

Piazza, M., Izard, V., Pinel, P., Le Bihan, D., & Dehaene, S. (2004). Tuning curves for approximate numerosity in the human intraparietal sulcus. *Neuron*, 44 (3), 547–555.

Piazza, M., Mechelli, A., Price, C.J., and Butterworth, B. (2006). Exact and approximate judgements of visual and auditory numerosity: an fMRI study. *Brain Res. 1106*, 177–188.

Piazza, M., Pinel, P., Le Bihan, D., and Dehaene, S. (2007). A magnitude code common to numerosities and number symbols in human intraparietal cortex. *Neuron 53*, 293–305.

Piazza, M., Fumarola, A., Chinello, A. and Melcher, S. (2011). Subitizing Reflects Visuo-spatial Object Individuation Capacity. *Cognition*, 121(1):147-153.

Piazza., M.(2010). Neurocognitive start-up tools for symbolic number representations. *Trends in Cognitive Sciences* 14(12):542--551, December 2010.

Pica, P., Lemer, C., Izard, V., & Dehaene, S. (2004). Exact and approximate arithmetic in an Amazonian indigene group. *Science*, 306, 499–503.

Pica, P., & Lecomte, A. (2008). Theoretical Implications of the Study of Numbers and Numerals in Munducucu. *Philosophical Psychology* 21:4, 507-522.

Pinel, P., Piazza, M., Le Bihan, D., and Dehaene, S. (2004). Distributed and overlapping cerebral representations of number, size, and lumi- nance during comparative judgments. *Neuron 41*, 983–993.

Pitt, B. & Casasanto, D. (2014). Experiential Origins of the Mental Number Line. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 36(36) 1174-1179.

Platt, J. R., & Johnson, D. M. (1971). Localization of position within a homogeneous behavior chain: Effects of error contingencies. *Learning and Motivation*, 2, 386–414.

Poincaré, H. (1910). Mathematical creation. *The Monist*, 20, 321

Porter, J.P. (1904). A preliminary study of the psychology of the English sparrow. *American Journal of Psychology* 15, 313-346.

Pöyhönen, Samuli. (2014). Explanatory power of extended cognition, *Philosophical Psychology*, 27:5, 735-759, DOI: 10.1080/09515089.2013.766789

Prior, A., Katz, M., Mahajna, I., & Rubinsten, O. (2015) Number word structure in first and second language influences arithmetic skills. *Frontiers of Psychology*. 6:266. doi: 10.3389/fpsyg.2015.00266

Pylyshyn, Z. W. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition* 32, 65–97.

Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition* 80(1/2), 127–158.

Pylyshyn, Z. (2003). *Seeing and visualizing: It's not what you think.* Cambridge, Ma: MIT Press.

Pylyshyn, Z. (2007). *Things and places: How the mind connects with the world.* Cambridge, MA: MIT Press.

Rav, Y. (2006). Philosophical problems of mathematics in the light of evolutionary epistemology. In R. Hersh (Ed.), *18 unconventional essays on the nature of mathematics* (pp.71–96). New York: Springer.

Railo, H., Koivisto, M., Revonsuo, A., Hannula, M.M. (2008) The role of attention in subitizing. *Cognition* 107: 82–104.

Regolin L, & Vallortigara G.(1995). Perception of partly occluded objects by young chicks. *Perception & Psychophysics* 57:971–976.

Relaford-Doyle, J., Núñez, R. E. (2018). Beyond Peano: Looking Into the Unnaturalness of Natural Numbers. In Bangu, S. (Ed). *The Cognitive Basis of*

Logico-Mathematical Knowledge. (pp.234-251). New York: Routledge

Renfrew, C., Morley, I. (2010). *The archaeology of measurement: comprehending heaven, earth and time in ancient societies*. Cambridge: Cambridge University Press.

Restle, F. (1970). Speed of adding and comparing numbers. *Journal of Experimental Psychology*, 83, 274–278. doi:10.1037/h0028573

Rey, G. (2014). Innate and learned: Carey, mad dog nativism, and the poverty of stimuli and analogies (yet again). *Mind & Language*, 29, 109–132.

Richardson, D. C., & Kirkham, N. (2004). Multi-modal events and moving loca- tions: Eye movements of adults and 6-month-olds reveal dynamic spatial indexing. *Journal of Experimental Psychology: General,* 133(1), 46–62.

Richerson, P., & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution.* Chicago: University of Chicago Press.

Rips, L. J., Bloomfield, A., and Asmuth, J. (2008a). From numerical concepts to concepts of number. *Behavioral and Brain Sciences*, 31, 623–642.

Rips, L. J., Asmuth, J. & Bloomfield, A. (2008b) Do children learn the integers by induction? *Cognition* 106:940–51.

Rivera, S.M., Reiss, A.L., Eckert, M.A., and Menon, V. (2005). Developmental changes in mental arithmetic: Evidence for increased functional specialization in the left inferior parietal cortex. *Cerebral Cortex 15*, 1779–1790.

Roberson, D. (2009). Color in mind, culture and language. In: Schaller, M., Norenzayan, A., Heine, S.H., Yamagishi, T. and Kameda, T. (Eds.) *Evolution, Culture and the Human Mind.* 167-184. Hove: Psychology Press.

Rolls, E.T. (2000). Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition. *Neuron 27*, 205– 218.

Ross, H. E., & Murray, D. J. (1996). *E. H. Weber on the tactile senses*. Hove, UK: Erlbaum, Taylor & Francis.

Rouse, J. (2007). Practice theory. In S. Turner & M. Risjord (Eds.), *Handbook of the philosophy of science: The philosophy of sociology and anthropology* (Vol. 15, pp. 499–540). Dordrecht: Elsevier.

Rugani, R., Vallortigara, G., Priftis, K., & Regolin, L. (2015). Number-space mapping in the newborn chick resembles humans' mental number line. *Science* 347, 534-536. doi:10.1126/science.aaa1379

Rupert, R.D. (2009). *Cognitive systems and the extended mind.* New York: OUP.

Sarnecka, B. W., Kamenskaya, V. G., Ogura, T., Yamana, Y., and Yudovina, J. B. (2007). From grammatical number to exact numbers: Early meanings of "one," "two," and "three," in English, Russian, and Japanese. *Cognitive Psychology*, 55(2), 136–168.

Sathian, K., Simon, T. J., Peterson, S., Patel, G. A., Hoffman, J. M., & Grafton, S. T. (1999). Neural evidence linking visual object enumeration and atten- tion. Journal of Cognitive Neuroscience, 11, 36–51.

Saxe, G. (1981). Body Parts as Numerals: A Developmental Analysis of Numeration among the Oksapmin in Papua New Guinea. *Child Development*, 52(1), 306-316. doi:10.2307/1129244

Skagenholt, M., Träff, U., Västfjäll, D., Skagerlund, K.(2018) Examining the Triple Code Model in numerical cognition: An fMRI study. *PLoS ONE* 13(6): e0199247. https://doi.org/10.1371/ journal.pone.0199247

Schlimm, D. (2018). Numbers through Numerals. In Bangu, S. (Ed). *The Cognitive Basis of Logico-Mathematical Knowledge*. (pp 195-217). New York: Routledge.

Schmandt-Besserat, D.(1996). *How Writing Came About*. Austin: University of Texas Press.

Schmandt-Besserat, D. (2010). The Token System of the Ancient Near East: its Role in Counting, Writing, the Economy and Cognition. In Renfrew, C. & Morley, I. (Eds.). pp.27-34, *The Archaeology of Measurement, Comprehending Heaven, Earth and Time in Ancient Societies*, Cambridge: Cambridge University Press.

Schlick, M. (1918), Allgemeine Erkenntnislehre,.Springer, Berlin. Excerpts in Kusch, M. (1995) *Psychologism: The Sociology of Philosophical Knowledge*. London ; New York: Routledge.

Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition* 80(1/2), 1-46.

Scholl, B. J. (2009). What have we learned about attention from multiple object tracking (and vice versa)? In D. Dedrick & L. Trick (Eds.), *Computation, cognition, and Pylyshyn* (pp. 49-78). Cambridge, MA: MIT Press.

Scholl, B. J., Pylyshyn, Z. W., & Franconeri, S. L. (1999). When are featural and spatiotemporal properties encoded as a result of attentional allocation? *Investigative Ophthalmology and Visual Science* 40(4), 4195.

Sears, C. R., and Pylyshyn, Z. W. (2000). Multiple object tracking and attentional processes. *Canadian Journal of Experimental Psychology*, 54(1), 1–14.

Seife, C. (2000). *Zero: the Biography of a Dangerous Idea.* New York: Viking.

Shaki, S., Fischer, M. H., & Göbel, S. M. (2012). Direction counts: a comparative study of spatially directional counting biases in cultures with different reading directions. Journal of Experimental Child Psychology, 112(2), 275–81.

Shea, N. (2011). New concepts can be learned: review of Susan Carey, The origins of concepts. *Biology and Philosophy*, 26, 129-139.

Siemann, J., & Petermann, F. (2018). Evaluation of the Triple Code Model of numerical processing—Reviewing past neuroimaging and clinical findings. *Research in Developmental Disabilities* 72 (2018) 106–117.https://doi.org/10.1016/j.ridd.2017.11.001

Simon, O., Mangin, J.F., Cohen, L., Le Bihan, D., and Dehaene, S. (2002). Topographical layout of hand, eye, calculation, and language-related areas in the human parietal lobe. *Neuron 33*, 475– 487.

Simon, O., Kherif, F., Flandin, G., Poline, J.B., Riviere, D., Mangin, J.F., Le Bihan, D., and Dehaene, S. (2004). Automatized clustering and functional geometry of human parietofrontal networks for language, space, and number. *Neuroimage 23*, 1192–1202.

Simon, T.J. (1997). Reconceptualizing the origins of number knowledge: A 'non-numerical' account. *Cognitive Development*, vol. 12 : 349-372.

Simon, T.J., Vaishnavi, S. (1996) Subitizing and counting depend on different attentional mechanisms: evidence from visual enumeration in afterimages. *Perception and Psychophysics* 58: 915–26.

Simons, D., & Levin, D.(1997). Change blindness. *Trends in Cognitive Sciences* 1, no. 7: 261–267.

Simons, D., & R. Rensink. 2005. Change blindness: Past, present and future. *Trends in Cognitive Sciences* 9, no. 1: 16–20.

Smith, C. L., Solomon, G. E. A., & Carey, S. (2005). Never getting to zero: Elementary school students' understanding of the infinite divisibility of number and matter. *Cognitive Psychology*, 51, 101-140.

Sober, E. (1978). Psychologism. *Journal for the Theory of Social Behaviour* 8: 165–92.

Soltész, F. & Szücs, D. (2014) Neural adaptation to non-symbolic number and visual shape: An electrophysiological study. *Biological Psychology* 103:203–11.doi: 10.1016/j.biopsycho.2014.09.006.

Spelke, E. S. Infants' intermodal perception of events. *Cognitive Psychology*, 1976, 8, 553-560.

Spelke, E. (2000). Core knowledge. *American Psychologist*, 55, 1233–1243.

Sperber, D. (1985). Anthropology and psychology: Towards an epidemiology of representations. *Man*, 20, 73–89.

Sperber, D. (1996). *Explaining Culture: a Naturalistic Approach.* Oxford, UK: Blackwell.

Sperber, D. (2006). Why a deep understanding of cultural evolution is incompatible with shallow psychology. in Nick Enfield and Stephen Levinson (eds.) *Roots of Human Sociality*, 431-449.

Starkey, P., & Cooper, R. (1980). Perception of numbers by human infants. *Science*, 210, 1033–1035.

Starkey, P., Spelke, E. S., & Gelman, R. (1990). Numerical abstraction by human infants. Cognition, 36, 97–128.

Sterelny, K. (2010). Minds: Extended or scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4),465-481. doi: 10.107/s1097-010-9174-y

Szűcs, D., & Csépe, V. (2005) The effect of numerical distance and stimulus probability on ERP components elicited by numerical incongruencies in mental addition. Brain Research. Cognitive Brain Research 22(2):289–300.

Tang, Y., Zhang, W., Chen, K., Feng, S., Ji, Y., Shen, J., Reiman, E., and Liu, Y. (2006). Arithmetic processing in the brain shaped by cultures. *Proceedings of the National Academy of Sciences of the United States of America*, 103:10775–10780.

Tieszen, R. (1990). Frege and Husserl on Number. *Ratio* 111: 150-164.

Trick, L.M., & Pylyshyn, Z.W. (1993) What enumeration studies can show us about spatial attention: evidence for limited capacity preattentive processing. *Journal of Experimental Psychology Human Perception & Performance* 19: 331–51.

Trick, L. M., & Pylyshyn, Z. W. (1994). Why are small and large numbers enumerated differently? A limited capacity preattentive stage in vision. *Psychological Review* 101 (1), 80-102.

Tudusciuc, O., and Nieder, A. (2007). Neuronal population coding of continuous and discrete quantity in the primate posterior parietal cor- tex. *Proc. Natl. Acad. Sci. USA 104*, 14513–14518.

Tzelgov, J., Henik, A., Sneg, R. & Baruch, O. (1996) Unintentional word reading via the phonological route: The Stroop effect in cross-script homophones. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22:336–49.

Uller, C., Hauser, M., & Carey, S. (2001). Spontaneous representation of number in cotton-top tamarins (Saguinus oedipus). *Journal of Comparative Psychology*, 115(3),248–257.

Valenza E, Leo I, Gava L, Simion F. (2006) Perceptual completion in newborn human infants. *Child Development* 77:1810–1821.

van Stigt, W. P. (1990). *Brouwer's intuitionism*, North-Holland, Amsterdam.

Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind*. Cambridge, MA: MIT Press.

Varley, R. A., Klessinger, N. J. C., Romanowski, C. A. J., & Siegal, M. (2005). Agrammatic but numerate. *Proceedings of the National Academy of Sciences of the United States of America*, 102, 3519–3524.

Verguts, T., & Van Opstal, F.2008. A colorful walk, but is it on the mental number line? Reply to Cohen Kadosh, Tzelgov, and Henik. *Cognition* 106(1):558-63.

Vetter, P., Butterworth, B., & Bahrami, B. (2008). Modulating attentional load affects numerosity estimation: evidence against a pre-attentive subitizing mechanism. *PloS one* 3 (9), e3269. http://dx.doi.org/10.1371/journal.pone.0003269.

Vetter P, Butterworth B, Bahrami B (2011) A candidate for the attentional bottleneck: Set-size specific modulation of right TPJ during attentive enumeration. *Journal of Cognitive Neuroscience* 23: 728–736.

Viarouge, A., Hubbard, E.M., McCandliss, B.D. (2014). The Cognitive Mechanisms of the SNARC Effect: An Individual Differences Approach. *PLoS ONE* 9(4): e95756. doi:10.1371/journal.pone.0095756

Wang, H. (1974). *From Mathematics to Philosophy*. London: Routledge & Kegan Paul.

Watson, D. G., Maylor, E. A., & Bruce, L. A. (2007). The role of eye movements in subitizing and counting. *J Exp Psychol Hum Percept Perform*, 33(6), 1389–1399.

Whalen, J., Gallistel, C. R., & Gelman, R. (1999). Nonverbal counting in humans: the psychophysics of number representation. *Psychological Science*, 10, 130–137.

Wheeler, M., Clark, A. (2008). Culture, embodiment and genes: unravelling the triple helix. *Philosophical Transactions of the Royal Society of London Series B* 363: 3563–3575.

Wiese, Heike. (2004). *Numbers, Language, and the Human Mind.* Cambridge ; New York: Cambridge University Press.

Williamson, L. L., Cheng, R.K, Etchegaray, M., & Meck, W. H. (2008) "Speed" Warps Time: Methamphetamine's Interactive Roles in Drug Abuse, Habit Formation, and the Biological Clocks of Circadian and Interval Timing. *Current Drug Abuse Reviews* 1 (2) 203-212(10).

Wilson, R. A. 1994. Wide computationalism. *Mind* 103: 351–372.

Wood, J. N., & Spelke, E. (2005). Infants' enumeration of actions: Numerical discrimination and its signature limits. *Developmental Science*, 8(2), 173–181.

Wood, G., Willmes, K., Nuerk, H-C., & Fischer, M. (2008). On the cognitive link between space and number: a meta-analysis of the SNARC effect. *Psychology Science Quarterly* 50(4): 489–525.

Wrangham, R. (2009). *Catching fire: How cooking made us human*. London: Profile.

Wu, W. (2014). *Attention*. London: Routledge.

Wynn, K. (1990). Children's understanding of counting. *Cognition*, 36, 155–193.

Wynn, K. (1992a). Addition and subtraction by human infants. *Nature*, 358, 749–750.

Wynn, K. (1992b). Children's acquisition of the number words and the counting system. *Cognitive Psychology*. 24 : 220–251.

Wynn, K. (1998). Psychological foundations of number: numerical competence in human infants. *Trends in Cognitive Sciences*, 2, 296–303.

Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition* 85(3),223–250

Xu, F., & Carey, S. (1996). Infants' metaphysics: the case of numerical identity. Cognitive Psychology, 30, 111–153.

Xu, F., & Carey, S. (2000). The emergence of kind concepts: A rejoinder to Needham and Baillargeon. *Cognition*, 74, 285–301.

Xu, F., & Spelke, E. S. (2000). Large number discrimination in 6-month old infants. Cognition, 74, B1–B11.

Xu, F. (2003). Numerosity discrimination in infants: evidence for two systems of representations. *Cognition* 89(1), B15–B25.

Xu, F., Spelke, E., & Goddard, S. (2005). Number sense in human infants. *Developmental Science*, 8(1), 88–101.

Xu, F., & Arriga, R. I. (2007). Number discrimination in 10-month-old infants. *British Journal of Developmental Psychology* 25:103–108.

Yantis, S., & Johnson, D. N. (1990). Mechanisms of attentional priority. *Journal of Experimental Psychology: Human Perception and Performance*, 16(4), 812–825.

Zebian, S.(2005). Linkages between number concepts, spatial thinking, and directionality of writing: The SNARC effect and the reverse SNARC effect in English and Arabic mono- literates, biliterates and illiterate Arabic speakers. *Journal of Cognition and Culture* 5(1–2), 165–90.