In Defense of Concept Variability

by

Alnica Visser

Bachelor of Arts, University of the Witwatersrand, 2011 Bachelor of Arts (Honours), University of the Witwatersrand, 2012 Master of Arts, University of the Witwatersrand, 2014

Submitted to the Graduate Faculty of the Dietrich School of Arts and Sciences in partial fulfillment of the requirements for the degree of Doctor of Philosophy

University of Pittsburgh

2023

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Alnica Visser

It was defended on

June 27, 2023

and approved by

Anil Gupta, Alan Ross Anderson Distinguished Professor of Philosophy, University of Pittsburgh

Kate Stanton, Assistant Professor of Philosophy, University of Pittsburgh

Dan Weiskopf, Professor of Philosophy, Georgia State University

Dissertation Director: James Shaw, Associate Professor and Department Chair of Philosophy, University of Pittsburgh Copyright © by Alnica Visser

2023

In Defense of Concept Variability

Alnica Visser, PhD

University of Pittsburgh, 2023

I present and defend the thesis of *Concept Variability*, the view that concepts can admit of variation in their representational contents without thereby losing their identity. I argue that the variability of concepts is central to their role in enabling cognition and thus that a concept's content variability is, despite philosophical orthodoxy to the contrary, a feature of cognition and not a bug.

I begin by arguing for the thesis negatively, by rejecting two prominent forms of *Concept Stability*, according to which concepts represent a stable set of representational contents. The first is *Criterial Stability*, according to which concepts represent the stable criteria of their own reference (e.g. criterial definitions or criterial essences). The second is *Type Stability*, according to which concepts represent stable types (e.g. prototypes or stereotypes). In each case, I argue that the *Stability* theses fail to capture the true scope of the sorts of contents that show up in how we use our everyday concepts. I close by providing two positive arguments for *Concept Variability*. The first is an abductive argument, according to which *Concept Variability* offers a better explanation than existing competitors for the range of empirical evidence showing persistence through content variability is a necessary condition for concepts to enable cognition, and thus retain their status as the building blocks of thought.

Table of Contents

Prefaceix
1.0 Introduction: Concept Stability & Concept Variability
2.0 The Problem(s) with <i>Criterial Stability</i>
2.1 Introduction7
2.2 Against <i>Definitionalism</i> 9
2.2.1 The Problem of Limited Cognitive Import11
2.2.2 The Problem of Typicality14
2.2.3 The Problem of Indeterminate Boundaries18
2.2.4 The Move to <i>Hybridism</i> 25
2.3 Against Hybridism25
2.3.1 The Problem of Cognitive Import, Again29
2.3.2 The Problems of Typicality and Indeterminate Boundaries, Again32
2.3.3 Prospects for <i>Hybridism</i>
2.4 Against Criterial Stability
3.0 The Problem(s) with <i>Type Stability</i>
3.1 Introduction
3.2 Statistical <i>Type Stability</i>
3.3 Non-Statistical Typicality 49
3.3.1 Ideal Concepts49
3.3.2 Dual Character Concepts53
3.3.3 Dependence Concepts55

3.3.4 Non-Statistical Typicality	58
3.4 Contextual Determinants of Typicality	59
3.4.1 Situations	59
3.4.2 Perspectives	61
3.4.3 Familiarity	62
3.4.4 Expertise	63
3.4.5 Contrasts	65
3.4.6 Contextual Typicality	66
3.5 Against Type Stability	66
4.0 Making Sense of Content Variation in Concept Use	69
4.1 Introduction	69
4.2 Preliminaries	. 73
4.3 Content Variation in Concept Use	. 75
4.3.1 Conceptual Development	76
4.3.2 Conceptual Sophistication	78
4.3.3 Inter-Disciplinary Differences	81
4.3.4 Intra-Disciplinary Differences	83
4.3.5 Cross-Cultural Differences	85
4.3.6 Cross-Contextual Differences	87
4.4 Against Invariantism	. 91
4.4.1 Invariantism	91
4.4.2 The Problem(s) with Invariantism	94
4.4.2.1 The Problem of Limited Cognitive Import Again, Again	. 95

4.4.2.2 The Problem of Empirical Invalidity	8
4.5 Against Contextualism	0
4.5.1 Contextualism10	0
4.5.2 The Problem(s) with <i>Contextualism</i> 103	3
4.6 Against Concept Stability 10	5
5.0 The Cognitive Necessity of <i>Concept Variability</i>	7
5.1 Introduction 10'	7
5.2 Concepts and their Many Contents 109	9
5.2.1 Concept Variability Articulated10	9
5.2.2 Concept Variability Exemplified11	3
5.2.2.1 Interim Objection 1: Don't Fish Constitute a Natural Kind? 116	6
5.3 The Enabling Role of Content Variation in Cognition	9
5.3.1 Categorization Requires Variable Activation of Epistemic Content124	0
5.3.2 <i>Cognition</i> Requires Variable Activation of Epistemic Content	2
5.3.3 Cognition Requires Variable Activation of <i>Representational</i> Content124	4
5.3.4 Interim Objection 2: Why not <i>Plural</i> Activation?124	6
5.3.5 Cognition Requires Variable <i>Storage</i> of Representational Content12	9
5.3.6 Interim Objection 3: Why Not Separate Storage?	1
5.4 Concept Variability is a Feature, Not a Bug13	6
6.0 Coda: Identity and Communication Without Stability 13'	7
6.1 On A Variable Concept's Identity13'	7
6.2 On Using a Variable Concept in Communication142	2
Bibliography	7

Table of Figures

Figure 1: Crayola Yellow-Greens	
Figure 2: "The Sandwich Alignment Chart"	
Figure 3: Various overlapping referents of FISH	

Preface

My first thanks go to the members of my committee, each of whom asked me hard questions while trusting me to come up with satisfactory answers: Anil Gupta, James Shaw, Dan Weiskopf, and especially Kate Stanton, who showed an unwavering enthusiasm for the project, and without whom I never would have had the courage to finish. Additional enormous thanks to my director and department chair, James Shaw, who served both as an intellectual and professional guide and a highly dependable bureaucratic support through an especially tumultuous time to live, study, and work many thousands of miles from my home.

I am equally grateful to many other faculty at Pitt: to Ken Manders, for pulling me aside on my first day as a graduate student at Pitt and asking me many difficult and probing questions about concepts and contents in the years that followed, many of which eventually led me down the path that led to this dissertation. And to Erica Shumener, for her expert guidance through my comprehensive examination and the subsequent revision and publication of the resulting paper, and for serving as the exemplar of a teacher and mentor I most hope to emulate throughout my philosophical career. Additional thanks to Bob Batterman, Tom Berry, Bob Brandom, Tom Ricketts, Nandi Theunissen, Michael Thompson, Lynne Tirrell, Japa Pallikkahayil, Jennifer Whiting, and Mark Wilson, each of whom formed a key part in my progress through the program and without whom my time at Pitt would have been significantly less elevating, edifying, and enjoyable.

Thanks also to all the wonderful staff both present and past: Jennifer Berkebile, Matt Ceruso, Collie Henderson, Connie Hrabovsky, Katie Labuda, Kathy Rivet, Delaney Szekely, and Diana Volkar. And from the international office: Caitlin Dawson and Richard Sherman. Each helped to smooth the rocky course through the various bureaucratic hurdles and burdens that can confront an international student at Pitt.

Above all, my heartful thanks to my peers, colleagues, and friends, all of whom have made my time in Pittsburgh happy and rich, and from whom I learnt more than I could ever have hoped. From my own beloved cohort, my dearest friends: Ash Purdy, Aaron Salomon, and Pablo Zendejas Medina. And to all the many esteemed and revered others: Mikio Akagi, Tom Breed, Ale Buccella, Will Conner, Laura Davis, Josh Eisenthal, Viv Feldblyum, Seth Goldwasser, Ned Howells-Whitaker, Dan Kaplan, Taylor Koles, Stephen Mackereth, Travis McKenna, Caleb Reidy, Jack Samuel, Alison Springle, Max Tegtmeyer, and Dan Webber. The same goes for my philosophical friends in other schools and departments: Osman Attah, Marina DeMarco, Eleanor Gordon-Smith, Kyle Landrum, Dejan Makovec, Mara McGuire, Siddharth Muthu Krishnan, Kyra Salomon, Jen Whyte, and Tom Wysocki. I have been lucky to know and learn from each of you.

Finally, to my family and husband, for whom it was never a question that I spend the better part of a decade on the other side of the world to read and write about interesting ideas and fun facts for no apparent reason other than the unwavering knowledge that I had to try. And to my unborn child, who I won't meet until a few months after this process is over, but who has helped to provide much needed motivation and perspective right at the end. I look forward to teaching you all about sandwiches, birds, fruit, fish, and more, and the very many different ways we all think about them.

1.0 Introduction: Concept Stability & Concept Variability

"Concepts are a mess" – (Murphy G., 2002, p. 492)

This is a dissertation about concepts. Except there is little consensus over what concepts *are*. (Margolis & Laurence, 2019) describe three options for the ontology of concepts: they are either (1) *mental representations*, or (2) *cognitive abilities*, or (3) *abstracta*. Of course, these options constitute not three *theories* of the nature of concepts, as such, but rather three *starting points*, for theorists regularly disagree over the metaphysics of mentality, representation, cognition, abilities, and abstractions, in addition to regularly disagreeing over how these three starting points might help to explicate any pre-theoretical notion of whatever we mean by the word "concept". Disagreement also exists over the whether these three starting positions must *exclude* one another or whether they might not be fruitfully combined. Indeed, as I see it, there is little reason to dismiss out of hand the idea that concepts exist as representations in the mind, while being realized in the expression of our cognitive abilities, in addition to being informatively typed and taxonomized by an appeal to abstracta.

My aim in this dissertation is not to arbitrate this debate, nor even to articulate any one version (or combination) of these metaphysical options. I mention them only in order to clarify my own starting point. I begin by assuming the first position, that concepts are mental representations. I do not hereby mean to commit myself to any particular representationalist theory of cognition, such as the *Language of Thought Hypothesis* (Fodor J. , 1975). Instead, I mean to commit myself only to a particular sort of *focus*.

The idea that concepts are to be understood as mental representations is the standard starting point in most contemporary work that approaches the study of mind and language *empirically*. It is the default position in both cognitive science and empirically informed philosophy of mind and language. This is the approach that asks us to make sense of human cognition and the expression thereof in thought, language, and action, by an appeal to the results sought and gathered in the empirical study of human brains, bodies, and behavior. I choose this starting point not because I think it is intrinsically superior to the many other sorts of approaches one might take with respect to the study of concepts. I choose it because I think the findings in the empirical literature are interesting, exciting, and hard to make sense of. And I choose it because I think I've found a new and compelling way to make sense of some of these findings.

A standard assumption shared across cognitive science, philosophy, and psychology is that the mental representations typically called "concepts" must represent at least some of their contents *invariantly*.¹ Concepts, it is said, must represent at least some *stable* set of their representational contents across the many various circumstances of their use, as well as across the many different cognizers that share in their use. I call this assumption *Concept Stability*. And my aim in this dissertation is to show that it is false. The representational contents of concepts can vary.

Some of the most exemplary articulations of *Concept Stability* can be seen in the following passages:

A central goal of cognitive science is to characterize the knowledge that underlies human intelligence. Many investigators have expended much effort toward this aim and in the process have proposed a variety of knowledge structures [i.e. concepts] as the

¹ Sometimes it's not "concept" but "category representation", "mental category", "category knowledge", "knowledge structure", "information structure" and the like. I will not concern myself with untangling these different usages here. These are all ways of referring to concepts as far as I'm concerned.

basic units of human knowledge, including definitions, prototypes, exemplars, frames, schemata, scripts, and mental models. An implicit assumption in much of this work is that knowledge structures are *stable*: Knowledge structures are stored in long-term memory as discrete and relatively *static* sets of information; they are retrieved intact when relevant to current processing; different members of a population use the *same* basic structures; and a given individual uses the *same* structures across contexts. These intuitions of *stability* are often compelling, and it is sometimes hard to imagine how we could communicate or perform other intelligent behaviors without *stable* knowledge structures. (Barsalou L., 1989, p. 76; my emphasis)

In experimental psychology, the distinctive aspects of mental life that enable each categorization are usually thought of as concepts. Shared mental structures are assumed to be *constant* across repeated categorizations of the same set of instances and different for other categorizations. (Keil F., 1994, p. 169; my emphasis)

[...] concepts are not essentially tied to any particular attitude: conceptual thought enables thinkers to enter into multiple mental states [...] with the represented content remaining *constant* across those attitudinal changes. [...] But for concepts to be context-independent in this sense, they must also be cross-contextually *stable*: the same concept must be able to be redeployed on different occasions and in different applications with a *common* representational import. (Camp, 2015, p. 593; my emphasis)

Keil and Camp state *Concept Stability* in order to endorse it. Barsalou states it in order to deny it. I am with Barsalou. On my view, concepts need not represent the same contents across the various contexts of their use, nor need they represent the same contents for cognizers who share a concept. Concepts can represent different contents on different occasions of their use. Different people might represent a great variety of different contents all while employing one and the same concept, and individual people may do the same across different contexts. There is no content that

must be represented across all (or even most) of a concept's use in cognition in order to count as an expression of that concept. I call this position *Concept Variability*.

Notice now that all three passages suggest an even stronger assumption than mere *Concept Stability*. Each author brings forth the idea that concepts must be stable *in order to enable cognition*. According to this stronger assumption, it is only *because* concepts represent (more or less) the same contents invariantly that it is possible for them to play a central role in our ability to think, perceive, and act in the many marvelous ways that we do. My own view is that cognition requires exactly the opposite. As I see it, cognition requires concepts that can vary. Not only *can* concepts admit variation in their representational contents, they *must* be capable of doing so if they are to enable cognition. By the end of this dissertation I will have defended both these claims.

The dissertation is structured as follows. In the first two Chapters I target two prominent forms of *Concept Stability*, showcasing how the assumption fails to make good sense of how concepts show up in everyday cognition, thus motivating the contrary idea that concepts can represent an unstable set of contents.

In Chapter 2, I target the view I call *Criterial Stability*. According to criterial accounts, concepts represent some stable set of criteria, which determine which objects belong among their referents. The classical version of such a view is *Definitionalism*, according to which concepts represent the definitions of their referents: the representational content of a concept is just an inventory of properties or characteristics that all and only its referents exhibit. The updated version of this view is *Essentialism*, according to which concepts represent those categories of objects that are at least presumed to share some more or less ineffable essence. In each case, I show that the account fails with respect to *scope*: not all concepts represent definitions and not all concepts represent referents presumed to share an essence. And I show that each account fails with respect

to *explanatory power*: neither definitions nor presumed essences provide adequate resources with which to make sense of how concepts are employed in the course of performing different cognitive acts.

In Chapter 3, I target the view I call *Type Stability*. According to type accounts, concepts represent stable types, representing that which is typical of a concept's referents. Such accounts score better than criterial accounts both in terms of *scope* and in terms of *explanatory power*. Nevertheless, these accounts still fail to adequately account fully for the large variety of concepts employed in everyday cognition, including the large variety of characteristics that can be represented by different sorts of concepts and the large variety of contexts that can influence the representational contents of concepts.

By the end of these two chapters, the idea that concepts can admit variation in their representational contents is firmly in view. In the final two chapters I defend this idea directly by arguing for *Concept Variability*, showing that concepts can, and indeed, must admit variation in their representational contents if they are to enable cognition.

In Chapter 4, I present an abductive argument, according to which *Concept Variability* offers a more satisfying explanation of some empirical evidence of content variation in cognitive behavior than those existing competitors that take seriously the idea that content might vary across different cognitive contexts. Here I target *Invariantism* and *Contextualism*, each of which aims to explain content variation while still retaining a commitment to some form of *Concept Stability*. I show here how the thesis of *Concept Variability* triumphs precisely where *Invariantism* and *Contextualism* falter.

I complete my defense of *Concept Variability* in Chapter 5 by defending the idea that concepts *must* admit variation in their representational contents if they are to enable cognition.

5

Assuming that concepts can be functionally defined as those cognitive entities that enable cognition, this idea allows me to offer a deductive argument for *Concept Variability*: concepts enable cognition, and they can do so only if they admit variation in their representational contents, hence concepts admit variation in their representational contents.

I close the dissertation with a look ahead, at some of the hard questions that open up once *Concept Variability* is accepted as well as a modest sketch of the answers I think worth exploring in future work in the empirical study of concepts and their role in cognition.

2.0 The Problem(s) with Criterial Stability

2.1 Introduction

Criterial Stability is the view that the representational content of a concept includes a *stable set of criteria that determine its reference*. The concept BIRD, for example, is understood as representing those criteria according to which birds, such as *chickens*, *robins*, *kiwis*, *ostriches*, and the rest, count as birds and therefore belong among the referents of BIRD.²

Criterial Stability is motivated primarily by the fact that it gives such a neat account of how concepts enable categorization. According to *Criterial Stability*, concepts represent the stable criteria of their own reference, and so they provide clear guidance in the categorization of particular objects: those that meet a concept's represented criteria of reference are among that concept's referents while those that don't, aren't. As (Margolis & Laurence, 1999) put it,

[...] something is judged to fall under a concept just in case it is judged to fall under the features that compose the concept. So, something might be categorized as falling under the concept CHAIR by noting that it has a seat, back, legs, and so on. Categorization on this model is basically a process of checking to see if the features that are part of a concept are satisfied by the item being categorized. (Margolis & Laurence, 1999, p. 11)

At its simplest, *Criterial Stability* holds that concepts represent nothing about their referents over and above that which determines their reference: BIRD represents about birds only

² I follow the convention of using CAPS to indicate reference to concepts themselves, rather than to any of their representational contents. I use *italics* to indicate reference to the (purported) contents of a concept. Thus "BIRD" refers to the concept of birds and "*robins*", "*ostriches*" and so on refer to the referential contents of BIRD.

that which makes various birds birds. The paradigm example of such a restrictive variant of *Criterial Stability* is *Definitionalism*, according to which the representational contents of concepts are exhausted by *stable definitions*. Such definitions are typically understood as representing a conjunction of individually necessary and jointly sufficient conditions for any object to be included among a concept's referents.³ BIRD, for example, is understood as representing only that which is individually and necessary and jointly sufficient for any object to be a bird; it represents, say, that birds are *feathered theropod dinosaurs*.⁴

Less restrictive versions of *Criterial Stability* allow that concepts can represent about their referents content over and above the stable criteria of their own reference. Such *Hybridism* allows that concepts can represent both criterial and non-criterial contents about their referents.⁵ *Essentialism*, for example, holds that concepts can represent stable criterial *essences* along with non-criterial *characteristics*, which are determined by the shared essence of a concept's referents and are thus typical among those referents.⁶ In this case, BIRD might represent both the criterial *feathered theropod dinosaur* and the non-criterial *winged flyer*, *nest builder*, *egg layer*, etc.

Restrictive versions of *Criterial Stability*, and especially *Definitionalism*, have long been out of favor in contemporary theories of concepts. And with good reason.⁷ It struggles to make

³ This is the view (Margolis & Laurence, 1999) call "The Classical Theory of Concepts." See also (Smith & Medin, 1981), (Murphy G., 2002) and (Margolis & Laurence, 2019) for helpful overviews as well as critical discussion in (Fodor, Garrett, Walker, & Parkes, 1980).

⁴ Here I use *italics* to refer to a concept's characterizing content, in addition to its referential content.

⁵ (Margolis & Laurence, 1999) refer to such Hybridist views as "Dual Theories." Examples include (Osherson & Smith, 1981), (Landau, 1982), (Erickson & Kruschke, 1998), and (Anderson & Betz, 2001). See also the critical surveys offered in (Weiskopf D., 2009) and (Machery & Seppälä, 2010). For hybridist responses, see (Gonnerman & Weinberg, 2010) and (Vicente & Martínez Manrique, 2016).

⁶ Also commonly called "Psychological Essentialism." Helpful surveys are offered by (Strevens, 2000) and (Neufeld, 2022). Different proponents of *Essentialism* vary on the scope of their view. Some, e.g. (Gelman S., 2003) restrict their account to concepts of natural kinds and some social kinds. More recent proponents, e.g. (Newman & Knobe, 2019) endorse a broader scope covering also concepts for a wider range of social kinds in addition to artifact kinds and individuals.

⁷ See, for example, (Smith & Medin, 1981), (Margolis & Laurence, 1999), (Margolis & Laurence, 2019), and (Murphy G., 2002).

sense of the large variety of concepts, their many apparent contents, and how these contents show up in everyday cognition. But some versions of *Hybridism* continue to hold sway in contemporary cognitive psychology. *Essentialism* especially continues to flourish as a particularly productive empirical hypothesis about the nature and role of kind concepts in higher cognition. The implication is that, in being less restrictive, *Hybridism* can overcome the problems that have motivated the demise of *Definitionalism*. Against this I argue here that no form of *Criterial Stability* can survive as a viable theory of concepts in general. That is, *concepts do not, in general, represent any stable criteria of their own reference*.

I begin in §2.2 by rehearsing three of the most important arguments against *Definitionalism*, which have motivated its well-known demise. In §2.3 I show how these arguments also undermine criterial versions of *Hybridism*, including *Essentialism*. I close in §2.4 by rejecting *Criterial Stability* in general.

2.2 Against Definitionalism

The paradigm of *Criterial Stability* is *Definitionalism*, according to which concepts represent only *definitions*. Most typically, a concept's definition is understood as articulating the individually necessary and jointly sufficient conditions for an object to be included among its referents. The definition represented by a concept thus serves as the criterion of its reference.

Apparently popularized by (Quine, 1961), the standard example of a definitional concept is BACHELOR, which can be defined by appeal to two conditions: being *unmarried* and being a *man*. On this definition all and only unmarried men are the referents of BACHELOR; something is a bachelor if and only if it is an unmarried man. The same may be said of other concepts in the

social domain, e.g. MOTHER may be defined by an appeal to the conjunction of being *female* and being a *parent*: all and only female parents are mothers, something is a mother if and only if it is a female parent. But concepts in radically different sorts of domains seem to admit of the same sort of treatment, e.g. TRIANGLE might be defined by appeal to being a *three-sided polygon*, SUNGLASSES by way of being *tinted eyeglasses*, CHEESE via being *pressed milk curds*, and, as I've already indicated above, BIRD with something like being a *feathered theropod dinosaur*.

For more complicated concepts, the represented definitions may not be nearly as easily articulated. Definitions once fully stated might be much more complex, and indeed in some particularly abstract domains, fully explicated definitions might seem permanently beyond our reach. Explicating the definitions for such concepts as KNOWLEDGE, JUSTICE, BEAUTY, and so on, makes for a familiar but ever-ongoing philosophical literature rooted in the methodology of *conceptual analysis*.⁸ Nevertheless, insofar as such analyses presuppose that our concepts already implicitly represent those definitions that determine their reference, waiting to be explicated by a sufficiently careful and thorough analysis, they are committed to some form of *Definitionalism* or other .⁹ For example, (Peacocke, 1998) as suggests:

Implicit conceptions involving definition may, though, be found in almost any domain. A significant segment of moral and political thought, for example, consists in making explicit the implicit conceptions and constraints which explain our applications of such notions as fairness, equality and opportunity. At the other end of the spectrum, I think we need to employ implicit conceptions in characterizing the mastery of even of some observational concepts. In mastering the concept *cube*, taken as an observational

⁸ For defense of such methodology, see for example (Jackson, 1998) and (Strevens, 2019).

⁹ Many contemporary conceptual analysts make no such presumption, and instead are committed only to the idea that correct definitions can be *developed* rather than be *discovered*. Such analysts engage in the methodology of *conceptual explication* in the tradition of (Carnap, 1950). These philosophers are thus not committed to the sort of *Definitionalism* I'm targeting here.

concept, a thinker must have an implicit conception with a content which includes this: that cubes are closed figures formed from square sides joined at right angles along their edges. Not all examples will be so trivial. In the case of any philosophically interesting concept, the question of the content of the implicit conception underlying it will be highly substantive. Answering the question will in such cases involve making substantive advance in the subject-matter.¹⁰ (Peacocke, 1998, p. 52)

Despite such suggestive examples and persistent methodological confidence, the objections against *Definitionalism* are plentiful. Three are especially noteworthy.¹¹ I will call them *The Problem of Limited Cognitive Import, The Problem of Typicality*, and *The Problem of Indeterminate Boundaries*. Each, I argue, is decisive against the plausibility of any *Definitionalist* account of concepts.

2.2.1 The Problem of Limited Cognitive Import

The core problem is this. By allowing the representation of only a conjunction of individually necessary conditions, *Definitionalism* predicts that a concept's content cannot include any content that concerns what is not common to all of its referents. And yet non-universal contents are indispensable in the successful employ of many ordinary concepts, especially for those categories of objects that appear to have no universally shared features. This has the highly problematic consequence that concepts can have little, if any, role to play in enabling higher cognitive behaviors. And yet concepts are ordinarily assumed to be the central components of

¹⁰ Peacocke notes that he doesn't assume *all* concepts to be explicated in this way. Some very simple concepts might admit of alternative treatments (perhaps, I conjecture, by way of *ostensive* definitions). As I'll make clear below however, I think the scope of *Definitionalism* is *much* narrower than such minor exception tolerance suggests.

¹¹ For a nice summary of a broader range of problems, see (Margolis & Laurence, 1999).

cognition, the building blocks of thought, those cognitive entities that function to enable higher cognition in general, including such processes as categorization, inference, communication, and the like. And so, if this presumed role of concepts is to be retained, they cannot be understood as representing definitions.

Consider again the case of BIRD. In order to be a definitional concept, it cannot include such features as the *ability to fly*, *having wings*, *having feathers*, *the ability to lay eggs*, or *having lightweight skeletons*, for none of these features are common to all birds. *Penguins*, *ostriches*, and other flightless birds can't fly, young chicks can't yet fly, and caged, clipped, diseased, or otherwise incapacitated birds are prevented from flying. *Moas* not only couldn't fly; they had no wings! Some victims of Psittacine beak and feather disease do not have feathers, and birds may be cleanly plucked without ceasing to be birds. Only reproductively mature birds can lay eggs, and only female and sufficiently healthy ones at that. And finally, some flightless bird species, the members of which have little need for light skeletons, have grown much denser bones than their flying cousins. And so, insofar as all these birds are indeed birds, none of these features can, by the lights of *Definitionalism*, be included in the definition that exhausts the content of BIRD.

In some particular extreme cases, the universal features that remain available for inclusion in a definitional concept are so general that they fail to distinguish objects that are among the referents of a concept from those that are not.¹² Again, the case may be made with BIRD. Without including the *ability to fly, having wings, having feathers, laying eggs*, and *having lightweight skeletons*, the concept can contain little more than the information that birds are *invertebrate*

¹² An early version of this concern is provided by (Wittgenstein, 2009) via "game", which, he argues, applies to a wide range of activities that bear merely a "family resemblance" to one another. On his account, there is no universal feature shared by all things falling in the range of the concept of a game, except perhaps the very general and uninformative feature of being an *activity*.

animals, which does not distinguish them from any of the other invertebrates. The prospects for defining them as *theropod dinosaurs* is equally dim.

Note then, that by allowing the representation of only content that concerns that which is universal to a concept's referents, *Definitionalism* renders many concepts all but cognitively inert. If only universal features are included in a concept's definition, but very few features characteristic of its referents are universal, very little can be communicated by or inferred from the concept alone about its referents. Characteristic but non-universal information cannot be conveyed by concepts that do not include them. Moreover, when the universal features are so general that they fail to distinguish those objects that are among a concept's referents from those that are not, definitional concepts can't even be used to perform categorizations. In either case, only once supplemented with additional information can definitional concepts become cognitively operational. On their own, concepts that represent only definitions have little role to play in higher cognition.

But the problem here goes deeper still. Requiring concepts to represent only definitions radically diminishes the potential scope of *Definitionalism*.¹³ For, if concepts can include only universal features, but there are cases where the universal features of a class of items together are not sufficient to distinguish between the members of the category from other items that also have those features—birds from invertebrate animals in general, or birds from theropod dinosaurs in general—and concepts must include conditions that are not only individually necessary but also jointly sufficient for a concept's reference to any object, then it follows that there are no concepts

¹³ The problem of scope has been presented elsewhere as *Plato's Problem*, namely the problem that we have found few if any plausible definitions for concepts. Definitionally concepts, if they exist at all, are very, very rare. See the discussion in (Murphy G. , 2002, pp. 17-19). But there is another problem of scope, brought forward by (Fodor, Garrett, Walker, & Parkes, 1980), who argue against reductive variants of *Definitionalism*. Such variants claim that it is only compound concepts that are definitionally structured, constituted by (comparatively) simpler concepts conjoined to form their definitions. The simplest concepts are used for building complex concepts, but aren't themselves complex, i.e. definitionally structured. And so, even if there are definitional concepts, such reductive variants of *Definitionalism*, are committed to the existence of non-definitional concepts.

for *internally variant categories*. Concepts, in other words, if they are to be definitionally structured, cannot exist for categories of objects that share no features that are at once common to all those objects while also being collectively distinctive of them. Absent an adequate definition for birds, and other similarly internally variant classes, *Definitionalism* thus implies that there exist no concepts for such categories!

Of course, this problem of scope also has consequences for the cognitive role of concepts: if there are no concepts for internally variant categories, then cognitive functions, like categorization, inference, and the like, which concern such categories must be performed without any concepts that concern the members of those categories in particular. It follows that no cognitive operations that concern birds, and other such categories can rely on concepts for them, if there are no possible definitions for them.

By requiring that concepts represent definitions then, *Definitionalism* is committed to viewing concepts as either of limited cognitive import, if a definition can be provided for a category of objects, or of no cognitive import at all, if no definition can be provided. In either case, if concepts are to play anything more than a peripheral role in cognition, they cannot represent definitions.

2.2.2 The Problem of Typicality

The basic phenomenon of interest here is that of *concept gradation*, known also as the phenomenon of *typicality effects*.¹⁴ Concepts across a wide variety of domains appear to admit of

¹⁴ Again see the helpful discussion in (Margolis & Laurence, 1999), (Margolis & Laurence, 2019), and (Murphy G., 2002). More recent findings are surveyed in (Machery E., 2009) and (Dieciuc & Folstein, 2018).

better and worse examples. The best examples are considered to be highly *typical* of a given concept's referents, while the worst are considered highly *atypical*. And such typicalities appear to influence a great variety of higher cognitive functions, including categorization, inference, learning, memory, and more. The problem for *Definitionalism* here is that definitions do not admit of such gradation in typicality: an object either meets the conditions of a definition or it does not. There is no discrimination between better or worse examples or more or less typical referents. The implication is that, on its own, an appeal to definitions cannot explain typicality effects. And so, if concepts represent only definitions, then some other aspect of cognition, beyond concepts, must be invoked to make sense of apparent concept gradation.

Consider again the primary example of a definitional concept, that of BACHELOR. Bachelors are unmarried men, but there are several other characteristics that are considered typical of unmarried men despite not being universal among them. The most typical bachelors, at least in contemporary Western imaginations, are single and young, perhaps in the early phases of their careers, perhaps still in college (likely pursuing a Bachelor's degree). They tend to either live alone (in a bachelor pad), or perhaps in a dorm or digs with other bachelors like them, in dwellings that typically include only rudimentary amenities. Further, bachelors are most typically not merely unmarried, but never married and they haven't fathered any children. Typical bachelors are eligible for marriage, and so are likely heterosexual—even in societies where gay men are now eligible for marriage, they often continue to be considered atypical grooms. And typical bachelors aren't in any serious or long-term romantic relationships; instead they go on dates and hook up, perhaps while still looking for "The One". In short, bachelors are the kind of guy that gets cast in *The Bachelor*.¹⁵

There are many ways for unmarried men to fall short of this bachelor type and so there are many sorts of atypical bachelor: widowers, divorcees, single fathers, boyfriends, fiancés, old bachelors, gay bachelors, asexual bachelors, priests, hermits, men in societies that don't recognize the institution of marriage, and more.¹⁶

The definitionally structured concept of bachelors cannot capture any of this nuance. First, because most of the information included in the bachelor type is not universal, the definitionally structured bachelor cannot include it. Only some bachelors conform to the bachelor type, and so the typical features cannot be included in the concept's definition. And second, because the type sorts between more and less typical bachelors, which a definition cannot do, discriminations between more and less typical bachelors cannot be informed by the definitionally structured BACHELOR. Concept definitions serve only to determine whether an object is among the referents of a concept or not; it doesn't contain any resources to rank or sort those referents.

And yet, there is a large and robust body of empirical evidence that shows that cognition depends on information about the typical features of category members. Most telling are results from *categorization* tasks. Cognizers are quicker to sort more typical category members than atypical members into their appropriate categories.¹⁷ (Rips, Shoben, & Smith, 1973) for example,

¹⁵ "5. All applicants must be single. To qualify as 'single', the applicant must not currently be involved in a committed intimate relationship, which includes: any marital relationship (whether or not the parties are separated or currently in the process of divorcing or annulling such marriage); any co-habitation relationship involving physical intimacy; or a monogamous dating relationship more than two (2) months in duration." "Eligibility Requirements", *Bachelor Nation*. Available at https://bachelornation.com/pages/eligibility/. Accessed June 8, 2020.

¹⁶ We may also note the possibility of married men who nevertheless conform to the bachelor type and may thus be classed as bachelors, of a sort. These "married bachelors", though married, live a "bachelor life", perhaps due to being separated from their spouse, whether romantically or domestically, or perhaps by eschewing various duties and norms that go with being a typical, or "good", husband.

¹⁷ See also (McCloskey & Glucksberg, 1978).

found that cognizers were quicker to categorize typical birds, such as (*North American*) robins and *blue jays*, as birds and slower to categorize less typical birds, like *chicken* and *geese*, as birds.

Cognizers are also quicker to produce more typical members as examples of a category rather than atypical members. In a study replicating and expanding on the findings of (Battig & Montague, 1969), (Van Overschelde, Rawson, & Dunlosky, 2004) for example, found that *eagles*, *robins*, and *blue jays* were listed as examples of birds more often (by a group of over 600 undergrads in three American colleges) than the much less typical birds, *chickens* and *vultures*.¹⁸

Typicality also affects *inference*. (Rips L., 1975) for example, found that participants were more likely to project newly learned information about the members of a sub-category to the category as a whole when the sub-category was of typical category members, and less likely to do so for sub-categories of atypical members. We are more likely to project information about *robins* to birds in general, than we are to project information about *vultures*. Similar findings are reported by (Osherson, Smith, Wilkie, López-Rousseau, & Shafir, 1990).

And finally, it has also been found that categories are learnt better, and faster, when cognizers are exposed to its more typical members than to its more atypical members.¹⁹ When acquiring competence with BIRD, we are better served by being exposed to *robins* and *ravens* than to *ostriches* and *penguins*.

Strictly speaking, these results aren't incompatible with *Definitionalism*. For one thing, the "mental categories" that participants in these studies are using to perform their cognitive tasks might not be concepts at all, in which case the fact that they display typicality effects has no bearing

¹⁸ See also (Mervis, Caitlin, & Rosch, 1976), (Barsalou L., 1983), and (Barsalou L., 1985).

¹⁹ At least for artificial categories: see, for example, (Posner & Keele, 1968), (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976), and (Mervis & Pani, 1980).

on the content of concepts.²⁰ But even if we do class these cognitive entities as concepts, it can still be that the concepts are structured definitionally, with typicality effects arising from some other aspect of our cognitive repertoire.

The trouble for *Definitionalism* is that it offers no resources with which to *explain* these typicality effects, and it certainly doesn't predict them. There is no reason we should expect a concept's referents to vary in their typicality with respect to the category to which the concept assigns them, if a concept functions merely to sort between its referents and its non-referents and contains no information about those features that are typical but non-universal for its referents. And there is no reason to expect that such typicalities would affect how cognizers go about learning and using concepts to categorize, infer, and communicate with respect to their referents, again because information about non-universal typicality is excluded from that concept. This means that positing a definitional content for concepts enable the cognitive production of categorization, inference, communication, and more. And so once again, the idea that concepts represent definitions is hard to square with the presumption that concepts are that which enable cognition in general.

2.2.3 The Problem of Indeterminate Boundaries

A related feature of concepts that places pressure on *Definitionalism* is that several sorts of concepts appear to admit of *indeterminate boundaries*. That is, at least some objects are neither

²⁰ Recall that some theorists deny that concepts are mental entities whatsoever, classing them instead as cognitive abilities, or as abstracta. See (Margolis & Laurence, 2019) for a helpful summary and bibliography of such views.

definitely included nor definitely excluded from the category of their referents. Definitions, by contrast, are usually thought of as having *sharp* boundaries; they allow no indeterminate cases. If an object meets the conditions of a definition, it is definitely in the category of objects for which it is true, but if it doesn't meet the definition, then it is definitely out of that category. Once again then, the idea that concepts represent only definitions cannot, on its own, then be invoked to explain the phenomenon of indeterminacy.

In the philosophical literature the most familiar example of such indeterminacies are color categories, which provide a common introduction to the phenomenon of *vagueness*.²¹ As the usual story goes, there are segments of the color spectrum that don't neatly fit into either of the color categories to which adjacent segments belong, which means that those color categories are *vague*. Consider, for example, the following range of Crayola colors (Fig. 1):



Figure 1: Crayola Yellow-Greens

²¹ See, for example, (Williamson, 1994) and (Keefe, 2000).

On the left is a color most would categorize as yellow. It is the color Crayola calls Brilliant Yellow, and the color Ford, Chrysler, and Humbrol call Yellow.²² Indeed, in the CMYK color space, #ffff00 corresponds to 100% Y, 0% C, M, and K; it is pure yellow. On the right is the color Crayola calls *Lime Green*, definitely green.²³ Between these two, things get less determinate. Consider the third color, which Crayola calls *Electric Lime*. It looks more green than yellow to my eye. Crayola agrees, classing it under its greens rather than its yellows.²⁴ And it is almost indistinguishable from the color Ford calls *Lime*. But it is also very close to what Chrysler calls *National Safety Yellow* and to the fluorescent yellows used for emergency vehicles and high visibility protective clothing in several parts of the world, including those worn by the *gilets jaunes*, the "yellow vest" protesters in 2018 in France.²⁵ Even trickier is the second color, which Crayola calls *Chartreuse*, also placed among its greens.²⁶ I find it almost indistinguishable from the yellow of #ffff00. Indeed, depending on the device or medium upon which the second square above is observed, it should be indistinguishable, for the web safe replacement for #ebff0d just is #ffff00, pure yellow. And so, although it is greener than pure yellow, Crayola's *Chartreuse* is nevertheless not green enough to merit a greener-than-yellow rendering in certain media. It is neither definitely yellow nor definitely green, and thus neither definitely in or nor definitely out of either color category. Both "yellow" and "green" are thus vague color terms, representing vague categories. And with respect to the concepts YELLOW and GREEN, both appear to be indeterminate in their referential boundaries, at the very least with respect to the Chartreuse.

²² <u>https://encycolorpedia.com/ffff00</u>, accessed July 6, 2020.

²³ <u>https://www.crayola.com/explore-colors/lime-green.aspx</u>, accessed July 6, 2020.

²⁴ <u>https://www.crayola.com/explore-colors/electric-lime.aspx</u>, accessed July 6, 2020.

²⁵ https://encycolorpedia.com/ebff0d, accessed July 6, 2020.

²⁶ <u>https://www.crayola.com/explore-colors/chartreuse.aspx</u>, accessed July 6, 2020.

An example more familiar to internet culture, and a growing number of philosophy undergraduates, is the infamous debate of the late 2010's over whether a hot dog is a sandwich. According to a viral tweet by @matttomic on May 1, 2017, there are nine positions one may take with respect to the boundaries of the sandwich category.²⁷ A handy chart (Fig. 2) summarizes:



Figure 2: "The Sandwich Alignment Chart"

According to the chart, "Hardline traditionalists" maintain that sandwiches must be composed of two pieces of bread (or similar baked product) along with a filling, sandwiched

²⁷ <u>https://twitter.com/matttomic/status/859117370455060481?s=20</u>, accessed July 6, 2020. See also this controversial chart: <u>https://www.reddit.com/r/funny/comments/8ba80v/the_sandwich_alignment_chart_we_deserve_oc/</u>, accessed July 6, 2020.

between the pieces, made of only the "classic sandwich toppings: meat, cheese, lettuce, condiments, etc.". On the opposite corner of the chart we find "Radical Sandwich Anarchy", a comically permissive position, allowing into the sandwich category any edible item composed of any kind of food enveloping in any way any kind of food. For the anarchist then, Pop-Tarts are sandwiches, as are empanadas, dumplings, folded pizzas, lasagna, samosas, calzones, filled crepes, sushi, cake, pie, pancake stacks, and of course, the bread sandwich, two or more adjacent slices of bread.²⁸

The internet controversy did not concern whether one ought to adopt either of these extreme positions; rather the bulk of the controversy concerned "True Neutral", whether a hot dog is a sandwich. Like a sub, the bread component of a hot dog is not divided into two separate pieces, and yet the bun is sliced so that it comes very close to operating like two pieces of bread, its parts being positioned both above and below the filling, holding it in place. And like a chip butty, a hot dog is filled with savory foods not that dissimilar from the traditional sandwich's meat, cheese, condiments, etc. Looking only at the fillings, there seems little that separates a ham and mustard sandwich from a frank and mustard dog. Nevertheless, the internet exploded with talk of hot dogs.²⁹ Indeed, the controversy is now so familiar among undergraduates of North American universities that it is a favorite example used by philosophy instructors to introduce their students to the distinction between necessary and sufficient conditions, and to how philosophers subject a

²⁹ Many of the most prominent and surreal twists and turns of the debate are covered by (Judkis, 2018), available at <u>https://www.washingtonpost.com/news/voraciously/wp/2018/12/12/a-hot-dog-is-a-taco-a-steak-is-a-salad-a-pop-tart-is-a-calzone-let-the-cube-rule-explain/</u>, accessed July 6, 2020. But see also the discussion from Merriam Webster: <u>https://www.merriam-webster.com/words-at-play/to-chew-on-10-kinds-of-sandwiches/hot-dog</u>, accessed July 6, 2020. And for a discussion of the tax implications of sandwich classification in New York State see (Abad-Santos, 2014), available at <u>https://www.vox.com/2014/7/22/5877325/new-york-law-sandwich-burrito-tax</u>, accessed July 6, 2020.

²⁸ Unsurprisingly, there exists an online quiz that will align you on the sandwich chart, on the basis of your judgments concerning particular cases. Take the quiz here: <u>https://isthisasandwich.netlify.app</u>, accessed July 6, 2020.

philosophically important concept or category to conceptual analysis. And the matter is not yet settled: as of July 2, 2020, after surveying 5366 US adults, <u>YouGov.com</u> found that 33% of respondents included hot dogs as sandwiches, 58% excluded them, and a full 10% didn't know how to class them.³⁰ With respect to SANDWICH then, it appears to admit of an indeterminate referential boundary.

These are not isolated instances. Indeterminate boundaries have been shown by empirical studies to be rampant in human cognition. And these cases exceed those that philosophers would group under the heading of vagueness. Paradigmatically, vague categories apply to objects that differ from one another along a *continuum*. Color is an exemplifying case, as are items with countable components, like heaps of sand and heads of hair—categories for which *Sorites* paradoxes can be constructed. But many categories can be indeterminate in their boundaries without being vague. Whether hot dogs are a kind of sandwich is a case in point. There is a determinate point at which a baked product is composed of two pieces rather than one; the question is whether the number of baked components can determine whether an edible product is a sandwich. And there is a determinate difference between a slice of ham and a frank, but it's not determinate difference between shelves and chairs, yet cognizers struggle to conclusively rule on whether shelves are furniture.³¹ And it's quite clear what differentiates chess from cricket, but

³⁰ <u>https://today.yougov.com/topics/food/survey-results/daily/2020/07/02/325f1/1</u>, accessed July 6, 2020.

³¹ For example, in a survey of 30 Princeton undergraduates, (McCloskey & Glucksberg, 1978) found a withinparticipant inconsistency rate of 0.4 on shelf-furniture judgments, i.e. a full 40% of their participants provided inconsistent judgements on whether shelves are furniture at different times. And yet, note that <u>ikea.com</u> lists its shelves and shelving units under its "Furniture" section (<u>https://www.ikea.com/us/en/cat/products-products/</u>, accessed July 7, 2020). With respect to another borderline type of furniture, curtains, about which the Princeton students displayed a 0.33 within-participant inconsistency rate is classed by Ikea under "Home Textiles" (again see <u>https://www.ikea.com/us/en/cat/products-products/</u>, accessed July 7, 2020).

it remains indeterminate whether chess is a sport.³² In these cases, the indeterminacy does not concern how much of some component an item must contain or manifest in order to be included in a category, but rather whether the presence or absence of some feature is sufficient for its inclusion in or exclusion from a category. Our working categories for sandwiches, furniture, sport, and more simply do not contain the resources to arbitrate on borderline cases, but the reason for this lack is not necessarily paradox-inducing vagueness.

If concepts represent definitions, it seems they should provide definite arbitration on borderline cases, like Crayola's *Chartreuse*, *hot dogs*, *shelves*, and *chess*.³³ That it can't be definitely determined which of these objects are and which aren't referents of YELLOW or GREEN, SANDWICH, FURNITURE, and SPORT, respectively, thus seems to indicate that the relevant cognitive entities by which we categorize them aren't definitional, or that we don't have a sufficient grasp of the definitions that provide such determinate delineation.

As with typicality effects, these boundary effects aren't strictly incompatible with *Definitionalism*. Perhaps these "mental categories" for which reference is indeterminate aren't concepts, and so their behavior in categorization carries no implication about the content of concepts. But then we encounter again the implication that those cognitive entities that are used to categorize objects aren't concepts, which suggests again that concepts are at best peripheral rather than central to human cognition. The alternative, that these mental categories are concepts, carries the same implication. If they are definitional after all, their definitions aren't what enable the

³² Consider: in 1999 The International Olympic Committee recognized chess as a sport and allowed an exhibition event at the Sydney Olympics before later denying the International Chess Federation's 2019 request to include chess in the 2024 Olympic Games. Nevertheless, chess has been included in recent iterations of the Asian Games, the South East Asian Games, and the African Games, all of which fall under the IOC's jurisdiction, in addition to the Pan-Armenian Games.

³³ Unless of course, the defining features included in a definition are themselves indeterminate. I'll return to this possibility below.

categorization of their referents. We successfully categorize a great many items as referents of YELLOW, GREEN, SANDIWCH, FURNITURE, SPORT, etc. and we draw good inferences and communicate successfully about categories that allow borderline cases despite having no access to, or without the existence of, determinately delineating definitions. Positing for concepts the representation of definitions can thus aid little in the project of theorizing concepts as those cognitive units that enable cognition.

2.2.4 The Move to Hybridism

All three of these problems seem to stem from the same fact: that *Definitionalism restricts* the content of concepts to the definitions that serve as the determinate criteria of their reference. Definitions represent too little content for concepts to play any significant role in our cognition. And definitions are of use in explaining why concepts appear to admit of neither typicality nor boundary effects. A natural response to such concerns is thus to *add more content*, content that can expand the potential cognitive import of concepts while also aiding in explaining the phenomena of typicality effects and indeterminate boundaries, without giving up on the idea that concepts represent the criteria of their own reference. This is the strategy of *Hybridism*.

2.3 Against Hybridism

The basic claim of any *Hybridist* account is that concepts represent a combination of contents of at least two different sorts. Of those that adhere to *Criterial Stability*, one of these sorts

must be *criterial* content, while the other can be *non-criterial*.³⁴ That is, in addition to content concerning that which is both individually necessary and jointly sufficient for a concept to refer to any particular object, such a criterial *Hybridism* holds that a concept can also represent *non-criterial contents*, contents that needn't serve as arbiters of a concept's reference and yet which can play a role in the acts of cognition that it serves to enable, possibly offering an explanation of the sorts of phenomena a bare appeal to criterial content cannot.

An early version of such *Hybridism* is the "core plus identification procedure" account proposed by (Osherson & Smith, 1981), according to which concepts represent both *definitions* and *types*, the former serving as ultimate criteria of reference ("the core") and the latter serving to enable individual instances of cognition, and especially categorizations ("the identification procedure"). A similar "symptoms and criterion" distinction is offered by (Landau, 1982), who holds that concepts represent both criterial *definitions* and symptomatic *types*, each of which is available for use in cognition. With respect to BACHELOR, for example, such *Hybridist* accounts might hold that the concept represents a definitional core, like *unmarried man*, alongside such noncriterial typicalities as *young, single, childless,* etc. Only the former determines which objects BACHELOR refers to, but both bodies of content are available for influencing and enabling particular BACHELOR-involving categorizations, inferences, communications, and so on.

Most prominent in the contemporary literature are various forms of *Essentialism*, according to which at least some concepts represent a combination of criterial *essences* and non-criterial *characteristics*.³⁵ The former is that which is represented as common among all referents of a

³⁴ Some *Hybridist* accounts, e.g. (Anderson & Betz, 2001), (Rice, 2016), (Vicente & Martínez Manrique, 2016), and (Reuter, 2019) posit multiple sorts of contents without requiring that one be criterial in particular. My interest here is only in those versions of that adhere also to *Criterial Stability*.

³⁵ Unlike most contemporary proponents of *Essentialism*, (Medin & Ortony, 1989) hold back from claiming that essences are *universal* to a concept's referents, allowing that some referents can fail to have the essence shared by
concept and which distinguishes them from all other objects. The latter are properties that are represented as being determined in some way or another by the shared essence of a concept's referents, but which need not be present among all those referents. The latter might be properties taken as being *caused* by underlying essences, or they may be taken as being *dependent on* or *correlated with* underlying essence in other ways.³⁶ At its most general, the essentialist idea is one of a sort of *connectedness* between criterial essence and typical characteristics:

People associate a concept with certain superficial features, but they do not regard those superficial features as sufficient for category membership; instead, they posit something further that unites them and explains how they are connected. This further thing allows them to answer the question: "What is it about these specific features that binds them together, so that it makes sense to associate all of them with this same concept?" (Newman & Knobe, 2019, p. 589)

And indeed, at its most general, the represented essence posited by *Essentialism* can be little more than that *something* in virtue of which the referents of a concept are all referents of that concept, and which is posited to lie behind the characteristics that form the rest of an essential concept's content. For example:

People may implicitly assume, for example, that there is some quality that bears have in common that confers category identity and causes their identifiable surface features, and they may use this belief to guide inductive inferences and produce explanations without being able to identify any feature or trait as the bear essence. This belief can be

their co-referents. They thus deny that the essence represented by a concept serves as absolute criterion for its reference, thereby denying *Criterial Stability*. Their view is therefore beyond the scope of my target here, but I will return to such a weakened idea in §2.4 below.

³⁶ Consider, for example, the teleological version of *Essentialism* offered by (Rose & Nichols, 2019) and (Rose & Nichols, 2020) and the domain general version in (Newman & Knobe, 2019).

considered an unarticulated heuristic rather than a detailed, well-worked-out theory. (Gelman S., 2003, p. 21)

With respect to BIRD then, for example, *Essentialism* would hold that the concept represents both what is essential to birds—whatever that might be—and that which is characteristic of them in virtue of being birds, which may include such exception-admitting typicalities as *feathered, winged flyer, nest-builder, egg-layer*, etc. And while both sorts of content can play a part in the cognitive acts a concept enables, it is only the essential part of a concept's content that serves as the criteria that ultimately determines its reference.

By adding a second sort of content to concepts, a criterial form of *Hybridism* seems well placed to address the problems facing more restrictive versions of *Criterial Stability*. For, the addition of typical characteristics appears to allow concepts to play a larger role in cognition while also offering the resources to explain the phenomena of typicality and indeterminate boundaries: unlike criterial definitions, typical characteristics needn't be universal among a concept's referents in order to be included in its content and they can admit of both typicality and borderline cases. That is, if hybrid, the concept BACHELOR can include both the definitional *unmarried man* and the merely typical *young single* and *childless* and can thus form part of the explanation of cognitions that rely on the latter. The additional of *winged flyer* and *nest-builder* in a hybrid BIRD can aid in explaining why *chickens* are judged to be less typical birds than *robins*, since it is only for the latter that these typicalities are true. The representation of the characteristic hue of *pure yellow* can explain why it isn't clear whether the merely somewhat yellow *Chartreuse* is a referent of hybrid YELLOW (or hybrid GREEN, for that matter). And the representation of the merely typical *two parts bread, one part savory filling* explains why it remains an open question whether

hot dogs are among the referents of a hybrid SANDWICH. Going hybrid, it seems, can protect *Criterial Stability* from the problems facing the more restrictive *Definitionalism*.

And indeed, it is true that the addition of typicality content can aid in according a wider role for concepts in cognition and in explaining how concepts admit of both typicality and indeterminate boundaries. But the problems facing *Definitionalism* cannot be overcome if such content is merely *added* to criterial content. By continuing to also posit criterial content, *Hybridism* retains the problems of *Criterial Stability*.

2.3.1 The Problem of Cognitive Import, Again

The upshot of *The Problem of Cognitive Import* is that, in allowing the representation of only universal content, a *Definitionalism* implies that concepts are either of minimal cognitive import, if definitional criteria can be given for a concept, or of no import at all, if no definitional criteria can be given. *Hybridism* relaxes this requirement by allowing the representation of non-universal content, thus allowing conceptual content to have much broader cognitive import.

And yet, precisely in virtue of this relaxation, the role of *criteria* in cognition remains minimal. For *Hybridism* allow that concepts can enable cognition without any appeal to criterial content. This is especially evident in the case of *Essentialism*, according to which competent cognizers need to know little, if anything at all, determinate about the essence presumed to unite the various referents of a concept in order to successfully employ that concept. That is, cognizers can be competent in a concept, applying it well in a great variety of cases, without being able to articulate anything about that essence they take to be what connects all the various referents of their concept.

This point is, in effect, what underlies the dispute between (Strevens, 2000), (Strevens, 2001), and (Ahn, et al., 2001). The latter contend that an appeal to essence is necessary in explaining certain cognitive behaviors, including both categorization and inference, especially in the cognition of children and in adults who harbor essentialist beliefs. Strevens counters that it is only category membership as such that influences cognition. Making the point in terms of children's inductions about the members of kinds, he claims:

What is important in accounting for the K-patterned projections [kind inductions] is that children believe that something about being a squirrel causes an animal to eat bugs. That the something is an essence adds nothing to the strength of the inference. (Strevens, 2000, p. 156)

And yet, as (Neufeld, 2022) notes on behalf of *Essentialism*, if the notion of essence is nothing more than *an indeterminate something*, then this dispute collapses: both parties agree that concepts can enable cognition without an appeal to any determinate essence. But it is only Strevens who sees the all-important consequence of this concession: there is no need to posit an essence among the represented content of concepts, or indeed any other determinate criterion of reference, in order to make sense of cognition, for cognition can proceed without the representation of such essential criteria.

Consider again, for example, the case of BIRD. And suppose that it represents nothing determinate about bird essence; it represents only that there is *something* in virtue of which all birds are birds along with such typicalities as the facts that birds *fly*, are *winged*, *feathered*, *lay eggs*, and so on. Given this concept and some object, we can reliably, though defeasibly, determine whether the object is a bird using only these typicalities. Similarly, given the concept and a bird, we can reliably, though defeasibly, infer that the bird can fly, has wings, lays eggs, etc. No appeal

to an indeterminate bird essence is necessary to perform such ordinary cognitions. The role of representing criterial essence in such cognitions is negligible.

The same is true of other forms of criterial *Hybridism*. As Armstrong, Gleitman, and Gleitman put it,

For some concepts, by hypothesis, there may be very little beyond the identification function [i.e. the type] that is stored in memory. For example, few, other than vintners and certain biologists, may have much in the way of a serious description [i.e. a definition] of *grape* mentally represented. (Armstrong, Gleitman, & Gleitman, 1983, p. 292)

And yet, there's no denying that a broad variety of cognizers, well beyond only "vintners and certain biologists," often and easily reason successfully with respect to GRAPE. The representation of criteria that determines its reference is therefore unnecessary for cognition to proceed.³⁷ It is enough to represent mere typicalities about its referents, such as their being *edible*, *seeded*, *grown on the vine*, in *bunches*, colored *green or red*, and *can be made into grape juice*, *vinegar*, *and wine*.

Note then, that although such *Hybridism* expands the potential content of concepts, it does not expand the cognitive import of *criteria* in particular. Indeed, it remains the case that for a large variety of concepts, no criterial content can be articulated by competent cognizers: recall especially the philosophically perplexing KNOWLEDGE, JUSTICE, BEAUTY, etc. And it remains the case that for several sorts of concepts, no absolute criterion of reference can be articulated, for there is nothing about the objects in internally variant categories that is at once common among all while

³⁷ For further discussion of this point, and the empirical data that motivates it, see, for example, (Hampton J. , 1979) and (Hampton J. , 1995).

being collectively distinctive of them: recall the case of BIRD. And so, for all these concepts, cognition must proceed without the representation of any criteria of reference. The role of criteria in cognition thus remains of limited, if any, cognitive import.

2.3.2 The Problems of Typicality and Indeterminate Boundaries, Again

The upshot in the case of *The Problem of Typicality* is that, by representing only universal criteria, *Definitionalism* cannot explain typicality effects, for typicalities aren't usually universal and universal criteria aren't typically graded. The addition of typicality information seems to present an easy fix. As Armstrong, Gleitman, and Gleitman put it, in terms of GRANDMOTHER:

[...] all it takes to be a grandmother is being a mother of a parent, but the difficulty is that all the same some grandmothers seem more grandmotherly than others. This issue is naturally handled in terms of a pair of representations: the first, the function [i.e. the type] that allows one to pick out likely grandmother candidates easily (it's probably that kindly grey haired lady dispensing the chicken soup) and the second, the description [i.e. the definition] that allows us to reason from *grandmother* to *female*. (Armstrong, Gleitman, & Gleitman, 1983, p. 293)

That is, by including the representation of types in the content of concepts, a criterial *Hybridism* can easily explain why some referents of a concept are judged to be better exemplars of that concept than others: it is because some referents of a concept better approximate a type represented by that concept than others.

And yet, once again, criteria do not play a role in this explanation. Although they can be of use in explaining another curious fact evident in some cases of concept gradation. Although the *typicality* of different concept referents is graded, their *membership* to the category of referents needn't be; both typical and atypical referents of a concept can be equally good *referents* despite not being equally good *exemplars* of a concept. And it is criteria of reference that are best able to explain this fact.

Compare, for example, different sorts of referents of BIRD: a *robin* and a *chicken*. The former is considered to be much more typical of BIRD than the latter. And part of what can explain this difference in typicality is that only the former has the ability to fly. And yet, chickens aren't therefore considered to be any less determinately birds. Both robins and chickens are both equal referents of bird: chickens are not like hot dogs. *Essentialism* especially has the resources to make sense of this phenomenon, by an appeal to essence: it is because chickens and robins share the bird essence that both are considered co-referents of bird, even though the former lack some of the most typical characteristics of birds.

And yet, precisely in virtue of this advantage, *Essentialism* cannot make sense of the *indeterminacy of referential boundaries*. And it certainly does not predict it. The role of essence is as an absolute criterion of reference. Proponents ordinarily take this to imply that essential concepts have "sharp category boundaries" (Neufeld, 2022, p. 2) with referents having "[a]bsolute category membership" (Gelman S. , 2003, p. 12). Indeed, several proponents of *Essentialism* take this feature to *distinguish* essentialist concepts from non-essentialist concepts:

[...] for natural kind categories, people should be more likely to give the endpoints of the scale when rating membership: a penguin is definitely a member of the bird category, even though it is not a very typical bird. For artifact categories, membership and typicality should be more closely coordinated. (Gelman S., 2003, p. 70)

Importantly, for *natural kinds*, no matter which degree of typicality was assigned to an animal, it was either judged to be a category member or not. In contrast, for *artifacts*,

category membership ratings are not independent of typicality. If an artifact is 'kind of' typical, it is also judged to be 'kind of' a category member. (Neufeld, 2022, p. 5)

That is, while essentialist concepts admit of gradation in typicality, they do not admit of indeterminate boundaries. Non-essentialist concepts, by contrast, like those that refer to artifact kinds, may admit of both.

Now, *Essentialism* does not usually purport to be a general view of concepts; their target domain is usually restricted to concepts of natural kinds, along with some combination of social kinds, functional kinds, and individuals, among others. There is thus room on such a hybridist account for non-criterial concepts. But other versions of criterial *Hybridism* tend to have a more general intended scope. And for these more general views, *The Problem of Indeterminate Boundaries* cannot be resolved by a mere addition of non-criterial content. If a concept includes criterial content, then we should expect cognizers to be able to employ it in order to resolve questions of an object's inclusion among the concept's referents. That they do not remains unexplained regardless of how much more non-criterial content is added to the representational content of a concept. And so, to the extent that some concepts admit of indeterminate boundaries, they shouldn't be understood as representing absolute criteria of their own reference.

2.3.3 Prospects for Hybridism

The problems for *Definitionalism* seems to stem from the restriction of conceptual content to mere criteria of reference. Criterial *Hybridism* addresses this by adding more content to concepts, typically by adding non-criterial typicalities to criterial definitions. In the case of *Essentialism*, criterial essence is combined with non-criterial characteristics. By such additions, some of the problems for *Criterial Stability* are softened. *The Problem of Typicality* seems especially well addressed by the addition of typicalities to the representational content of concepts. Nevertheless, by retaining the representation of criteria, *Hybridist* accounts continue to fall short of solving *The Problem of Cognitive Import* and *The Problem of Indeterminate Boundaries*, for criteria continue to be of limited cognitive import, where they can be represented at all, and they continue to fail to explain the phenomenon of indeterminate boundaries. The true problem is positing in concepts the representation of criteria for their own reference. The problem is *Criterial Stability*.

2.4 Against Criterial Stability

Going hybrid is not the only option available to proponents of *Criterial Stability*. Alternative options include expanding the operative notion of a *definition*, according to which definitions can be any Boolean combination of conditions, the whole of which serves as criterion of reference and yet no part of which need be a necessary condition of reference. Consider, for example, this *disjunctive* definition for baseball's FAIR BALL:

In baseball, a batted ball is a fair ball if and only if it settles on fair ground between home and first base or between home and third base, or is on or over fair territory when bounding to the outfield past first and third base, or touches first, second, or third base, or first falls on fair territory on or beyond first base or third base, or, while on or over fair territory, touches the person of an umpire or player. (Machery E. , 2011, p. 17) In this case, none of the listed conditions are *necessary* for a batted ball to be fair, but each is sufficient. None of the represented content is therefore universal among the concept's referents, and yet the combination of disjuncts provides a necessary and sufficient criterion of reference. This move thus easily treats internally variant classes, but the prospects for dealing with typicality effects and indeterminate boundaries are less certain: such a definition offers no indication of typicality, nor indeterminacy of reference.

Another alternative is relaxing the requirement that criteria of reference be *universal* among a concept's referents. Medin and Ortony, for example, suggest that essences can be represented as merely typical among a concept's referents:

[...] it may be part of the represented essence of *bird* that birds fly, even if it happens that not all birds do fly and that people know this. (Medin & Ortony, 1989, p. 184)

Such a relaxation would have the benefit of accommodating both internally variant classes and indeterminate boundaries, but only by giving up on the idea that represented criteria serve as absolute arbiters of reference. Essence becomes merely another non-criterial characteristic of a concept's referents.

Another option is to recast criteria as themselves *indeterminate*. This option is typically pursued by applying the results of fuzzy logic to the psychology of concepts.³⁸ Zadeh motivates the idea as follows:

More often than not, the classes of objects encountered in the real physical world do not have precisely defined criteria of membership. For example, the class of animals clearly

³⁸ See, for example, discussion of these themes in (Belohlavek & Klir, 2011).

includes dogs, horses, birds, etc. as its members, and clearly excludes objects as rocks, fluids, plants, etc. However, such objects as starfish, bacteria, etc. have an ambiguous status with respect to the class of animals. [...] such imprecisely defined "classes" play an important role in human thinking [...] The purpose of this note is to explore in a preliminary way some of the basic properties and implications of a concept which may be of use in dealing with "classes" of the type cited above. The concept in question is that of a *fuzzy set*, that is a "class" with a continuum of grades of membership. (Zadeh, 1965, p. 338)

Such an amendment would thus understand concepts as admitting, in general, of indeterminate boundaries, while also making room for some internally variant classes, but again only by giving up the idea that concepts represent absolute criteria of their own reference.

Each of these options offers a potentially viable amendment to the idea that concepts represent about their referents the criteria of their reference, but only by retaining a problematic notion of criteria or by turning criteria into contents that look awfully similar to the non-criterial types and characteristics posited by *Hybridism*. They are contents the parts of which can be non-universal, merely typical of a concept's referents, and which need not serve as ultimate arbiters of reference. A natural suggestion then, is to give up on *Criterial Stability* and endorse in its stead a *Typicalist* account, according to which concepts represent about their referents only *types*, i.e. typicalities characteristic of their referents. It is this suggestion to which I turn my attention in the next Chapter.

3.0 The Problem(s) with *Type Stability*

3.1 Introduction

Type Stability is the view that concepts include among their representational contents only *stable types*. On this view, concepts represent only the *stable characteristics* typical of their referents, none of which need be universal among them.³⁹ This account thus allows that BIRD, for example, can represent such characteristics as *winged*, *feathered*, *egg-layer*, *nest-builder singer*, and *flyer*, among others, even though none are universal among birds.⁴⁰ And sandwich can represent *edible*, *savory filling*, *enveloped in two pieces of bread*, and so on, despite the fact that sandwiches don't universally exhibit such characteristics.

The primary motivation for *Type Stability* is its natural interpretation of *typicality effects*, i.e. the fact that different referents of a concept are judged by cognizers to be more or less typical of that concept than others. *Ostriches* and *chickens*, for example, aren't judged to be nearly as typical of BIRD as are *robins* and *ravens*. And *chip butties* and *subs* are judged to be quite a bit more atypical of sandwich than are *BLTs* and *PB&Js*. The idea that concepts represent the typical characteristic *flyer* that *robins* are judged to be better examples of birds than *chickens* and it is because sandwich represents *enveloped by two pieces of bread* that *subs* are

³⁹ As will become clear below, I make no claim that characteristics *must* be non-universal among a concept's referents, only that they can be.

⁴⁰ I continue the convention of using small caps for reference to concepts themselves and italics for their contents, including the characteristics they represent and their (candidate) referents.

judged to be worse examples than *PB&Js*. It is the typicality content of concepts themselves that are enough to explain such observed typicality effects.

Some views, like *Hybridism*, which I discussed in the previous Chapter, can offer a similar explanation of typicality effects, since it admits among the cognitive contents of concepts both typicality content and criterial content. But *Type Stability* also offers the resources to explain several other observed facts about concepts, which have spelled trouble for more traditional accounts, including the different variants of *Criterial Stability*, according to which concepts represent the universal criteria of their own reference.

First, insofar as typical characteristics aren't necessarily universal among the referents of a concept, type concepts can represent a broader range of contents than can criterial concepts. This expansion allows concepts to play a wider role in the cognitions that implicate or depend upon non-criterial content. The fact that our bird-related cognitions, for example, frequently depend on such non-criterial contents as *egg-layer* and *nest-builder*, characteristics that are not universal among birds, is directly explicable from the fact that the type concept BIRD represents them. And the fact that our sandwich-related cognitions often draw on non-criterial contents like *savory filling* and *enveloped by two pieces of bread* is immediately explicable from the fact that the type concept SANDWICH represents them. No further concept-external contents or processes are necessary to explain the possibility of such cognitions.

Second, insofar as typicality itself is gradable, the phenomenon of *indeterminate boundaries* can also be explained. Some concepts admit of borderline referents, objects that are judged as neither determinately included among a concept's referents nor determinately excluded from them. It is, for example, indeterminate whether *hot dogs* are included among the referents of SANDWICH and it is indeterminate whether *chartreuse* is included among the referents of

39

YELLOW. *Type Stability* readily explains this phenomenon too: different objects can approximate to a greater or lesser extent the typicalities represented by a concept and so, to the extent that a concept's reference is determined by how closely an object approximates that content, different objects can be more or less determinately included among that concept's referents.⁴¹ Moreover, such an explanation isn't blocked as it is in the case of *Criterial Stability*, since type concepts, unlike criterial concepts, can fail to issue determinate verdicts on their reference, for they needn't represent any criterial characteristics.

Different versions of *Type Stability* are to be distinguished primarily by how they explicate the murky notion of *typicality*. The earliest and most prominent version of the view is *Prototypicalism*, according to which typicality is to be understood as a matter of *statistical frequency*.⁴² On this view, concepts function to keep track of the *stable statistical regularities* present among their referents. In particular, the idea is that characteristics are typical of a concept's referents to the extent that they are at once most commonly present among referents while also being most commonly absent among non-referents. As Rosch and Mervis introduce the idea:

⁴¹ Some early proponents of *Type Stability* conflated typicality effects and boundary effects such that atypical referents are always also borderline referents. But more contemporary accounts have been careful to separate the two phenomena. Objects can be atypical referents without also being borderline referents: *penguins* and *chickens*, for example are atypical referents of bird, but they aren't also borderline referents. *Pterodactyls*, by contrast, could be counted as both atypical and borderline, insofar as they, much like the most prototypical birds, were *flying dinosaurs*. ⁴² An early statement of *Prototypicalism* is given by (Rosch & Mervis, 1975) with support from their accompanying empirical results reported in (Heider, 1972), (Rosch, 1973a), (Rosch, 1973b), (Rosch, 1975a), (Rosch, 1975b), (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976), (Rosch, 1977), (Rosch, 1978), and (Mervis & Rosch, 1981). Inspired by this work, there is now a wide range of prototypicalist accounts in the literature; see especially (Tversky, 1977), (Hampton J., 1979), (Hampton J., 1995), and (Hampton J., 2006). For a helpful orientation to the main contours of the view and the bodies of evidence commonly cited in support of it, see (Smith & Medin, 1981), (Margolis & Laurence, 1999), (Murphy G., 2002), and (Margolis & Laurence, 2019). And for a more recent survey of some important results and debates in the contemporary literature, see (Dieciuc & Folstein, 2018). See also the recent literature on "conceptual spaces", e.g. (Douven & Gärdenfors, 2019), which makes essential use of the Roschian notion of a prototype.

The more prototypical a category member, the more attributes it has in common with other members of the category and the less attributes in common with contrasting categories. (Rosch & Mervis, 1975, p. 602)

Prototype concepts thus prioritize for representation those characteristics of their referents that exhibit both *high in-category frequency* and *low out-category frequency*.⁴³ Taking their cue from Wittgenstein, Rosch and Mervis call such typicality, "family resemblance".⁴⁴ Others call it "central tendency" or "cue validity." I will call it *prototypicality*. On this version of *Type Stability* then, concepts represent *stable prototypes*, representing only that which is stably *prototypical* of their referents.

One alternative elaboration of typicality is *stereotypicality*. *Prototypicalism* emphasizes those characteristics that are *in fact* most common and distinctive of a concept's referents. But what cognizers *believe* to be most prototypical of a concept's referents can come apart from actual prototypicality, insofar as cognizers' statistical knowledge (or beliefs) about a concept's referents and their characteristics is incomplete, incorrect, or otherwise imperfect or impaired. Such *Stereotypicalism* builds this possibility into the notion of typicality itself by holding that concepts represent only that which cognizers *believe* to be prototypical of their referents. Call such believed

⁴³ Too often proponents of *Prototypicalism* are explicit only about the high in-group frequency aspect of prototypicality. But, as (Murphy G., 2002) makes clear, low out-group frequency is just as critical, and likely implicit in much of *Prototypicalism*. If concepts were to represent what is most common of their referents regardless of its distinctiveness, they would have to represent an enormous amount of information about what is most common to all existents, purely in virtue of their referents being existents. Moreover, empirical investigations of categorization have shown that it is the *distinctiveness* of conceptually represented characteristics, rather than their *commonality*, that is of central importance in categorization. Commonality, meanwhile, is of central importance in inference; see especially the results surveyed in (Markman & Ross, 2003). And so, if concepts are to represent content that is at once useful for categorization and for inference, while also minimizing useless redundancies, they must prioritize the representation of content that is at once common and distinctive of their referents.

⁴⁴ See especially (Rosch & Mervis, 1975) and (Wittgenstein, 2009). But note that Rosch and colleagues should not be read as scholars of Wittgenstein's notion of family resemblance. They are offering only an empirically informed explication of his idea.

prototypicality *stereotypicality*. On this view, concepts represent only that which is stably *stereotypical* of their referents.⁴⁵

Despite these epistemic differences, both versions of *Type Stability* are committed to the idea that typicality is *statistical*, and thus that the typical characteristics selected for representation by concepts are determined by what is (at least believed to be) their stable in-category and out-category *frequencies*. On both accounts then, concepts represent only statistical contents, which thus exhausts the sort of contents concepts can contribute to the processes of cognition.

Similarly committed to the exclusive cognitive import of statistical contents is *Exemplarism*.⁴⁶ Distinctive of *Exemplarism* is that concepts are analyzed as representations of the set of their known exemplars, with cognitions being determined by the extent to which newly encountered objects approximate the characteristics exhibited most frequently by previously encountered exemplars. *Prototypicalism* and *Stereotypicalism*, by contrast, analyze concepts as *summary representations* of the characteristics of a concept's referents, rather than those referents to a summary of the characteristics exhibited most prototypically (or stereotypically) by a concept's referents, rather than to the characteristics of each known exemplar in particular. And

⁴⁵ Although many proponents of *Type Stability* tend to declare allegiance to *Prototypicalism*, most are better understood as committed to some version of *Stereotypicalism* insofar as they rely on their experimental participants' own judgments and reports of the prototypicalities of the referents of their concepts. Many of these proponents also tend to conflate *Prototypicalism* and *Stereotypicalism* insofar as they rely on subjective judgments of prototypicality in studies of real categories (i.e. those for which participants already have a concept, e.g. BIRD and SANDWICH), but on objectively measured prototypicality in studies of artificial categories (i.e. those for which participants must learn a concept in an experimental context).

⁴⁶ See, for example, (Medin & Schaffer, 1978), (Hintzman & Ludlam, 1980), (Hintzman, 1986), (Brooks, 1987), and especially the work of Nosofsky and colleagues: (Nosofsky R., 1984), (Nosofsky R., 1986), (Nosofsky R., 1988), (Nosofksy, 1991), (Nosofsky & Johansen, 2000), (Nosofsky & Zaki, Exemplar and Prototype Models Revisited: Response Strategies, Selective Attention, and Stimulus Generalization, 2002), (Palmeri & Nosofsky, 2010), and (Nosofsky, B., & Kumar, 2020). See also the overviews provided by (Margolis & Laurence, 1999), (Margolis & Laurence, 2019), and (Murphy G., 2002), as well as the latter's more recent skepticism over whether *Exemplarism* constitutes an account of concepts distinct from *Prototypicalism* and *Stereotypicalism* at all, in (Murphy G., 2016).

yet, the views are united in treating concepts as contributing to cognition only statistical contents about their referents insofar as each view treats concepts as contributing to cognition only those characteristics that are (at least believed to be) statistically most frequent among their (known) referents and statistically least frequent among their (known) non-referents.

My aim in this Chapter is to show that *Type Stability* as much as *Criterial Stability*, fails to capture the true scope of conceptual contents. The idea that concepts represent typical characteristics of their referents is hard to deny in the face of robust and overwhelming empirical evidence about the effects of typicality on cognition. But the idea that typicality is always, or even frequently, a matter of *frequency* is quite a different matter. And indeed, the idea that statistical typicalities are represented *stably* becomes clearly untenable once we consider the realities of everyday cognition.

I begin in §3.2 with a more detailed explication of *Type Stability* and its commitment to a statistical interpretation of typicality. In §3.3 I present a series of examples of different types of concepts, each of which I argue are best understood as representing *non-statistical* typical characteristics of their referents. In §3.4 I bolster my argument against statistical interpretations of typicality by reviewing the evidence for the idea that typicality can be determined by *qualitative* factors, such as *context*, rather by than only *quantitative* factors, such as *statistical frequency*. I close in §3.5 by reviewing how the evidence for non-statistical typicality and context-sensitivity undermines statistical interpretations of typicality in particular as well as *Type Stability* in general.

3.2 Statistical Type Stability

Some versions of *Type Stability* have a *flat* conception of concept typicality, according to which concepts simply represent the typical characteristics of their referents as being *on a par* with one another. On such views, an object counts as a referent of a concept just in case it meets a sufficient number of any combination of the represented typicalities. As Margolis and Laurence put this idea in terms of BIRD:

[...] if BIRD is composed of such features as FLIES, SINGS, NESTS IN TREES, LAYS EGGS, and so on, then on the Prototype Theory, robins are in the extension of BIRD because they tend to have all of the corresponding properties: robins fly, they lay eggs, etc. However, BIRD also applies to ostriches because even though ostriches don't have all of these properties, *they have enough of them*. (Margolis & Laurence, 1999, pp. 27-28; my emphasis)

On such an account, the referents of concepts are understood as more or less typical of that concept to the extent that they exhibit more or less characteristics represented as typical by that concept:

Generally speaking, exemplars that [...] *share a higher number of features* will also have higher goodness-of-example ratings. (Dieciuc & Folstein, 2018, p. 492; my emphasis)

For example, as Murphy presents the findings reported by Rosch and Mervis,

They [i.e. Rosch and Mervis] found that the five most typical examples of furniture (chair, sofa, table, dresser, and desk) had thirteen attributes in common. In contrast, the

five least typical examples (clock, picture, closet, vase, and telephone) had only two attributes in common. For fruit, the five most typical items had sixteen attributes in common, but the least typical ones had absolutely no attributes in common. (Murphy G. , 2002, p. 34)

Other versions of *Type Stability* hold that the typicalities represented by concepts are *weighted*, ranked as *more* or *less* typical than one another.⁴⁷ That is, a concept such as BIRD can rank its typical characteristics as more or less typical of birds and thus as more or less important for those cognitions about birds, such as categorizations and inferences, enabled by the concept. Being *winged*, or *feathered*, for example, may be more heavily weighted than *flyer*, since the latter is less typical of birds than the former, but both may still be represented by BIRD insofar as each is typical of birds. Again Murphy offers a helpful example:

The concept is represented as features that are usually found in the category members [i.e. the referents of the concept], but *some features are more important than others*. It is important for weapons that they be able to hurt you, but not so important that they be made of metal, even though many weapons are. Thus, the feature "can do harm" would be highly weighted in the representation, whereas the feature "made of metal" would not be. (Murphy G. , 2002, pp. 43-44, my emphasis)

On such versions of the account, concept reference is determined by a slightly more complicated route. Objects still count as referents of a concept to the extent that they meet a sufficient number of typical characteristics, but now such sufficiency is influenced by the weight of each characteristic. It is, in general, better for a candidate referent to meet a few highly typical

⁴⁷ See, for example, discussion in (Tversky, 1977) and (Smith, Osherson, Rips, & Keane, 1988).

characteristics than it is to meet several less typical characteristics. As Murphy explicates such a process of categorization:

After going through the object's features, one adds up all the weights of the present features and subtracts all the weights of its features that are not part of the category. If that number is above some critical value, the *categorization criterion*, the item is judged to be in the category; if not, it is not. Thus, it is important to have the highest weighted features of a category in order to be categorized. For example, an animal that eats meat, wears a collar, and is a pet might possibly be a dog, because these are all features associated with dogs, though not the most highly weighted features. If this creature does not have the shape or head of a dog, does not bark, does not drool, and does not have other highly weighted dog features, one would not categorize it as a dog, even though it wears a collar and eats meat. So, the more highly weighted features an item has, the more likely it is to be identified as a category member. (Murphy G. , 2002, pp. 43-44; my emphasis)

One advantage of employing a weighted conception of typicality is that it accommodates better than the flat notion the sorts of examples and intuitions that recommend various versions of *Criterial Stability* without implicating *Type Stability* in the same sorts of trouble. The idea that concepts represent criterial characteristics, e.g. defining or essential characteristics, can be accommodated by *Type Stability* if concepts are understood as weighting such characteristics most highly for their referents. Being *unmarried* and *male* can be weighted more heavily than *single* and *looking* in the concept BACHELOR and being *justified*, *believed*, and *true* can be ranked as the most typical characteristics of the referents of KNOWLEDGE. This can then explain why such characteristics seem most important in the categorization of different types of men (in the case of BACHELOR) and beliefs (in the case of KNOWLEDGE), without thereby also requiring that concepts, in general, represent absolute criteria of their own reference. Instead, the general idea is

simply that concepts represent that which is typical of their referents and those characteristics that are most typical are proportionately most important for categorizing those referents.

Weighted typicality also allows for more nuanced explanations of *typicality effects*. Different exemplars of a concept can exhibit the same number of characteristics typical of a concept's referents but still differ in their overall typicality ranking to the extent that the characteristics of some exemplars are weighted as more typical than others. It's possible, for example, to distinguish between the typicality of *vultures* and *ostriches* with respect to BIRD. Each sort of exemplar lacks one important characteristic typical of birds: ostriches don't fly and vultures don't build nests.⁴⁸ But to the extent that *flyer* is weighted more heavily than *nest-builder* for BIRD, we can predict that *vultures* will be judged to be more typical than *ostriches*. On the flat interpretation of concept typicality, by contrast, *vultures* and *ostriches* would be deemed equally typical birds, equally good exemplars of bird.

Additional nuances are also available for explaining *boundary effects*. Weighted typicality allows the prediction that a concept will admit of borderline referents insofar as it represents no characteristics weighted so highly as to be effectively criterial or near-criterial. BACHELOR, for example, has a rather determinate boundary insofar as *unmarried* and *man* are both weighted so highly as to be virtually criterial. SANDWICH, by contrast, represents none such highly weighted contents, and thus should admit of a less determinate boundary.

Weighted typicality also allows the identification of *which* referents of a concept are its borderline referents: they are the candidate referents that exhibit medium overall typicality—not high enough to be determinately in the concept's reference, but not low enough to be determinately out—due to, for example, representing typical characteristics of only medium weight, or only a

⁴⁸ Both also share the atypicalities of being large and being unable to sing.

medium number of typicalities of high weight. An example of the latter case: we can explain why it's indeterminate whether *hot dogs* are included among the referents of SANDWICH by appeal to the fact that they exhibit two highly weighted characteristics, *edible* and *savory filling*, but are missing another, *enveloped by two pieces of bread*. And an example of the former: we can explain why *chartreuse* is on the boundary of YELLOW insofar as its focal hue is a medium distance to, rather than close to, the highly typical *focal yellow* of YELLOW.⁴⁹

The idea that concepts represent typicalities—and especially weighted typicalities—thus seems a rather promising general account of concepts. And yet throughout the literature committed to different forms of *Type Stability* persists the additional assumption that typicality, along with corresponding typicality weightings, are determined by *frequency*. That is, *flyer* is highly typical of bird insofar as more birds fly and fewer non-birds don't, but *nester* is less typical insofar as fewer birds nest and more non-birds do. Similarly, the idea is that *savory filling* is highly typical of sandwich in virtue of being very frequent among sandwiches and very infrequent among non-sandwiches.⁵⁰

But, as might be already evident from these examples, it's not so obvious that it is such stable frequencies, correctly believed or not, that always lie behind typicality. And indeed, I

⁴⁹ This explanation of why *chartreuse* is an indeterminate referent of YELLOW might be rather quick for some. Here's one way to fill out the details: the focal hue of chartreuse is in between the focal hues represented by YELLOW and GREEN. The mostly highly weighted characteristic of the referents of YELLOW is a *close approximation to focal yellow*. And the most highly weighted characteristic of the referents of GREEN is a *close approximation to focal green*. But the focal hue of *chartreuse* isn't close enough to the focal hue of either yellow or green to be determinately in the reference of either YELLOW or GREEN. Instead its hue is a medium distance away from each focal hue, and is thus indeterminately a referent of each concept.

⁵⁰ In the case of *Prototypicalism*, the idea is that such stable frequencies are *in fact* the case for concepts, like BIRD, SANDWICH, and the rest. For *Stereotypicalism*, the claim is that these are these stable frequencies at least *believed* to be the case. And for *Exemplarism*, the claim is that these are the stable frequencies that are the case, or perhaps merely believed to be the case, for *known exemplars*.

contend that it is quite frequently *non-statistical* factors that determine which characteristics concepts represented as most typical of their referents.

3.3 Non-Statistical Typicality

Not all concepts represent *statistical* typicalities. Some represent typicalities that depart from what would be included in the prototype (or stereotype) of a concept's referents (or known referents). These non-statistical typicalities are characteristics represented as typical of a concept's referents, resulting in the familiar sorts of typicality effects and boundary effects, without also being represented as statistically most common and distinctive of those referents. The typicality of these characteristics is thus rooted in something other than their (at least believed) statistical typicalities thus serve as counterexamples to the statistical interpretations of typicality implicit in the statistical variants of *Type Stability*: concepts do not represent only the statistical typicalities.

3.3.1 Ideal Concepts

A particularly illustrative example of non-statistical typicality is found in *ideal concepts*.⁵¹ These are concepts that weight most highly those characteristics that their most ideal referents

⁵¹ See, for example, (Barsalou L., 1983), (Barsalou L., 1985), (Medin, Lynch, Coley, & Atran, 1997), (Lynch, Coley, & Medin, 2000), (Bailenson, Shum, Atran, Medin, & Coley, 2002), (Medin & Atran, The Native Mind: Biological Categorization and Reasoning in Development and Across Cultures, 2004), (Burnett, Medin, Ross, & Blok, 2005), (Voorspoels, Vanpaemel, & Storms, 2011), and (Voorspoels, Storms, & Vanpaemel, Idealness and Similarity in Goal-Derived Categories: A Computational Examination, 2013). On *goal-directed* concepts, which represent characteristics

exhibit, or would exhibit if circumstances were right. Given sufficiently non-ideal circumstances, however, the actual referents of an ideal concept will fall short of many or even all the ideal characteristics represented by such a concept. This means that what is actually, or at least believed to be, most statistically common and distinctive of a concept's referents can diverge from what is considered most ideal for them. Nevertheless, it is the represented ideal characteristics rather than (or in addition to) the statistically frequent characteristics of an ideal concept's referents that underly its observed typicality and boundary effects. It is the most ideal referents of a concept that are judged to be its best examples and judged to be furthest away from its referential boundary. And so, we have strong reason to include among the typicalities represented by concepts *ideal typicalities*, i.e. characteristics that are judged most typical of an ideal concept's referents without also being statistically common and distinctive of those referents.

Before we get to some examples, note the basic difference between ideal concepts and the concepts posited by statistical variants of *Type Stability*. The statistical interpretation of typicality predicts that the best examples of a concept are its *most average* referents, those referents that most closely approximate the characteristics that are represented as most common and most distinctive of a concept's referents. But a good example of an ideal concept is one which best approximates the ideal characteristics of a concept's referents, regardless of their prevalence among those referents. What is ideal, that is, might be quite different from what is prototypical or stereotypical of a concept's referents. But, insofar as typicality and boundary effects are determined by ideal characteristics, rather than, or in addition to, statistically frequent characteristics, the notion of typicality must, at least in the case of ideal concepts, admit of non-statistical typicality.

ideal with respect to a contributing to the achievement of a goal, see (Barsalou L. , 1985), (Borkenau, 1990), (Davis & Love, 2010), and (Rein, Goldwater, & Markman, 2010).

The most commonly cited example of such an ideal concept, first discussed by (Lakoff, 1987), is that of DIET FOOD, whose referents are ideally, but rather infrequently, *zero-calorie food*. Almost no actual diet food can fully approximate this ideal (except *shirataki noodles* and perhaps *water*, if it counts as a food). Diet foods are much more commonly *low-calorie* food. This is also what is most distinctive of them. And so, if typicality is statistical, then the most typical and most central exemplars of diet food must be low-calorie food; both medium to high-calorie food and very low to zero-calorie food should be judged *less* typical and *more* borderline, in proportion to their distance from the prototypical low-calorie food. And yet, intuitively, diets foods are more typical and more central the lower their caloric content; their typicality and centrality should be *directly proportional* to their caloric content. That is, the very best examples—those judged to be most typical—and those judged to be furthest away from the concept's referential boundary, should be those that best approximate the (almost) unobtainable extreme of having no caloric content.

Empirical investigations have confirmed such intuitions in several different cases. (Lynch, Coley, & Medin, 2000), for example, found that tree experts judge the typicality of trees in proportion to their *height* and *weediness*; the tallest and least weedy referents of TREE are judged to be its best examples and furthest away from its referential boundary. (Atran, 1999) found that the best and most central examples of BIRD according to Itza Mayan participants were those that were most *meaty* and *delicious*.⁵² Burnett and colleagues report that fishers rank the typicality of local referents of FISH according to their local cultural *desirability*, most often those that offer "good eating". (Burnett, Medin, Ross, & Blok, 2005, p. 5) And (Foster-Hanson & Rhodes, 2019)

⁵² See also (Medin & Atran, 1999).

found that children tend to judge the typicality of ANIMAL by ideality even where their adult counterparts do not, e.g. they judge the typicality of skunks in proportion to their *stinkiness*.

Similar findings are reported by (Horstmann, 2002) with respect to various facial expressions: the more intensely an emotion is expressed, the more typical of that emotion the expression is judged. Consider, for example, the findings in Study 2 for anger and surprise. In each case the typicality ratings (represented by filled dots) correspond closely to the ideality ratings (represented by filled diamonds), both of which increase with expressive intensity, and depart significantly from the statistical frequency ratings (represented by unfilled diamonds), which are higher at medium intensity and lower at both low and high intensities.⁵³ That is, it appears that typicality judgments can correspond to characteristic *extremes* rather than statistical *averages*, strongly suggesting that typicality can be non-statistical.

In each of these cases, both typicality judgements and boundary judgments are influenced by the approximation of a candidate referent to one extreme of a spectrum: diet food is best to the extent that its caloric content is *least*, trees are best to the extent that they are *most* tall and *least* weedy, the best birds are the *most* meaty and the *most* delicious, the best fish are the *most* culturally desirable, offering the *best* eating, the best skunks are the *most* stinky, and the best facial expressions are those that exhibit the *most* emotive intensity. But if typicality were always a matter of statistical frequency, the best diet foods, trees, birds, fish, skunks, and facial expressions would be those of *medium* caloric content, height, deliciousness, desirability, stinkiness, emotive

⁵³ The unfilled dots represent familiarity ratings, which also depart from statistical frequencies, and thus suggest another form of non-statistical typicality. I will discuss exemplar familiarity in §3.4.3 below.

intensity, etc. And so, to the extent that not all typicalities are those characteristics that are most common and most distinctive among a concept's referents, not all typicality is statistical.⁵⁴

3.3.2 Dual Character Concepts

A closely related sort of concept, which can diverge similarly from the concepts posited by statistical variants of *Type Stability* is *dual character concepts*, which are disjunctive multi-modal concepts that represent two independent characterizations of their referents.⁵⁵ One characterization is *descriptive*, representing the actual (or believed to be actual) characteristics of a concept's referents. The other is *normative*, representing the characteristics participants judge referents *should* have, regardless of how many referents actually fulfill such norms.⁵⁶ These concepts are *disjunctive* in that they admit among their referents objects that meet only the descriptive characterization, as well as those that meet both.

⁵⁴ Note, however, that it is not the representation of extreme values, as such, that is incompatible with statistical typicality. See, for example, discussion in (Levering & Kurtz, 2006), (Palmeri & Nosofsky, 2010), and (Kim & Murphy, 2011). For, in some cases, approximation to extreme values may be highly distinctive of a concept's referents. Consider, for example the concept PERSONAL BEST, which has as its referents the very best attempt an athlete has achieved in their sport. Ideally, these attempts are the *fastest, furthest, heaviest, highest* etc. efforts of an athlete's sporting history. But insofar as these attempts remain among the referents of PERSONAL BEST only to the extent that they are not exceeded by a new attempt, being the fastest/furthest/heaviest/highest and so on is also highly common and distinctive of such attempts.

⁵⁵ See, for example, (Knobe, Prasada, & Newman, 2013), (Del Pinal & Reuter, 2015), (Leslie, "Hillary Clinton is the Only Man in the Obama Administration": Dual Character Concepts, Generics and Gender, 2015), (Del Pinal & Reuter, 2017), (Phillips, De Freitas, Mott, Gruber, & Knobe, 2017), and (Reuter, 2019). See also the closely related philosophical notion of a *thick concept*, e.g. (Kirchin, 2013), which are concepts that conjoin interdependent descriptive contents and evaluative contents, but which has received almost no empirical investigation, except a preliminary study by (Willemson & Reuter, 2021). ⁵⁶ Alternatively, the normative characterization represents what referents *would* have, under more normatively ideal

⁵⁶ Alternatively, the normative characterization represents what referents *would* have, under more normatively ideal circumstances. See, for example, discussion in (Del Pinal & Reuter, 2015) and (Leslie, "Hillary Clinton is the Only Man in the Obama Administration": Dual Character Concepts, Generics and Gender, 2015).

The dual character concept ARTIST, for example, represents both descriptive characterizations, e.g. *produces artworks* and *produces art professionally*, and normative characterizations, e.g. *produces works of great aesthetic value* and *artistically talented*.⁵⁷ And these characterizations are independent, because a person can count as a referent of ARTIST on the basis of meeting only the concept's descriptive characterization, or by meeting only the normative, or by meeting both: one can be an artist by making superficial art for a living, or by making art of great aesthetic value recreationally, or by being a professional master.

The problem for the statistical interpretation of typicality implicit in *Type Stability* is that the normative contents of a dual character concept, much like the idealities represented by an ideal concept, can come apart from that which is most common and distinctive of its referents. And yet, it has been found that the best and most central examples of dual character concepts are those that meet its *normative* characterization, thus strongly suggesting that dual character concepts include among their representation contents non-statistical *normative typicalities*.

Producing works of great aesthetic value may be distinctive of artists, but it's quite uncommon among them. In fact, the same is true for being a professional artist: while producing art professionally is distinctive of artists, it's not at all common among artists in general. The great majority of artists create their art in their spare time and their work is, on average, of average aesthetic value. This is also most distinctive of artists: non-artists don't produce much (if any) art, professionally or recreationally, regardless of quality. The most statistically common and distinctive characteristics of the referents of ARTIST then, if it represents artists in general, are *produces works of middling aesthetic value* and *produces works recreationally*. According to the statistical interpretation of typicality, *Type Stability* thus predicts that the best and most central

⁵⁷ I follow here the findings and characterizations of ARTIST reported by (Knobe, Prasada, & Newman, 2013).

examples of artist are hobbyists of rather average talent. And yet, it is talented masters and professional successes—Banksy, Kahlo, Warhol, O'Keeffe—that are judged better and more central examples of artist. The concept must therefore represent at least some non-statistical *normatively typicalities*, in addition to more or less idealized descriptive typicalities.

The same findings have been reported for various other social role concepts, e.g. SCIENTIST, TEACHER, MOTHER, and FRIEND (Del Pinal & Reuter, 2015), as well as JAZZ (Reuter, 2019), HAPPINESS (Phillips, De Freitas, Mott, Gruber, & Knobe, 2017), ROCK MUSIC and LOVE (Knobe, Prasada, & Newman, 2013). In each of these cases, much as in the case of ideal concepts, we find a concept that represents some (perhaps unobtainable) abstract value—a normative ideal—of which actual referents can and often do fall significantly short. Because of this, what is statistically most common and distinctive of their referents can, at least in some cases, come apart from what is judged to be most typical of them.⁵⁸ And so, to the extent that some normative typicalities aren't statistically typical of a concept's referents, the typicality represented by concepts cannot be exhausted by statistical typicality.

3.3.3 Dependence Concepts

Lastly, consider the case of *essential concepts*, or what may be dubbed more broadly, *dependence concepts*, already discussed at length in the previous Chapter under the heading of *Essentialism*. These are concepts that represent not only the typical characteristics of their referents, but also the casual (or metaphysical) dependencies that exist between those

⁵⁸ Once again one must be careful not to collapse normative ideality and statistical extremes. Some statistical extremes can be sufficiently common and distinctive of a concept's referents to be included in its prototype or stereotype. The claim here is only that there are *some* normative typicalities that are statistical extremes that aren't included among that which is (or believed to be) statistically common and distinctive of a concept's referents.

characteristics. Insofar as statistical analyses of typicality treat the properties concepts represent as independent variables, (believed) prototypes must exclude any dependencies represented by dependence concepts.⁵⁹

In essential concepts, which are usually applied to the members of kinds (or categories thought to be kinds), referents are represented as sharing some underlying essential characteristics upon which their more superficial characteristics depend.⁶⁰ The essence is maximally common and maximally distinct, for all and only the members of the same kind share the same essence. To this extent the statistical interpretations of typicality can accommodate essential concepts. But it is the representation of the dependencies holding between an essence and the other typical characteristics of a concept's referents, as well as the dependence relations that hold between those characteristics, that spells trouble for the statistical approach to typicality.

Consider again the example of BIRD. The typical properties prevalent among birds seem rather well captured by the statistical approach: birds typically *fly*, have *wings*, have *hollow bones*, *live in nests*, and so on. What the approach fails to capture however, is the extent to which these features *co-vary*. It is only those birds that can fly that typically live in nests. And it is only because they have wings and hollow bones that flying and nesting are even possible for them.

The same goes for FRUIT: fruits are typically *sweet*, *fleshy*, *edible*, and so on. But it is the fleshy and sweet fruits that are also usually the edible ones. And it is certainly only these that go well in fruit salads; it is edibility along with a fleshy texture and sweet flavor upon which the property of going well in a fruit salad depends. Most crucially, it is the fact that fruit typically

⁵⁹ See, for example, the discussion and results reported in (Malt & Smith, 1984), (Murphy & Medin, 1985), and (Barsalou L. , 1993).

 $^{^{60}}$ See, for example, discussion in (Ahn, et al., 2001), (Ahn, Marsh, Luhmann, & Lee, 2002) , (Gelman S. , 2003), (Gelman S. , 2004), (Leslie, 2013), and (Neufeld, 2022).

envelop seeds that forms a crucial part of the explanation of their sweetness and fleshy texture, as these are the properties that are most appetizing to the animals who are most likely to eat and spread the seeds to new fertile grounds.

And yet, on the basic "family resemblance" analysis of concepts, the prevalence of each property present among a concept's referents is assessed independently.⁶¹ The statistical analysis for each property is conducted in isolation from all others. The (believed) prototype is the summary of these statistical results: only those properties that are (believed to be) at once most common and most distinctive of a concept's referents that are represented in their (believed) prototype. But dependence concepts represent more than the average prevalence of each characteristic in isolation; they represent the dependencies that exist between typical characteristics. And so, to the extent that dependence concepts represent information that goes beyond what may be represented by the (believed) prototype for their referents, those concepts do not represent only that which is statistically typical of those referents.

As before, the importance of non-statistical content comes out most clearly in the influence of dependencies on typicality and boundary judgements. Several studies have found that the typicality of the referents of dependence concepts are determined not according to how prototypical or stereotypical their properties are but rather according to how central they are to the dependence relations of importance in the category.⁶² These studies consistently find that children

⁶¹ *Type Stability* can be modified to include the representation of feature correlations; see, for example, (Smith & Medin, 1981). But see also some reservations for the feasibility of such an account in (Murphy G. , 2002). The subsequent literature tends to agree with Murphy. Most theorists who find evidence of dependency content in concepts have used it to motivate the rejection of *Type Stability* in favor of *Essentialism* or some other "Causal Theory" of concepts, often referred to in the literature as "Theory Theories" or "Knowledge Theories" of concepts. See, for example, reviews in (Margolis & Laurence, 1999), (Murphy G. , 2002), (Machery E. , 2009), and (Margolis & Laurence, Concepts, 2019).

⁶² See, for example, (Gelman & Markman, 1986), (Gelman & Markman, 1987), (Keil F., 1989), (Rips L., 1989), (Gelman S., 2003), and (Rehder, 2003). See also the review of findings in (Ahn & Kim, 2001) and (Neufeld, 2022).

and adults will prioritize those characteristics considered to be essential or otherwise explanatorily central to the properties of a concept's referents in general over those properties that are considered to be more accidental, superficial, or otherwise explanatorily peripheral to them, regardless of the relative prevalence or absence among referents and non-referents. (Gelman & Markman, 1986) and (Gelman & Markman, 1987) found that participants will continue to categorize a leaf insect as an insect even if shares many more superficial (and especially perceptual) properties with a leaf, conforming better to the leaf prototype than the insect prototype. And (Keil F. , 1989) found that participants will continue to categorize an animal as a skunk, even if it has been artificially altered to conform better to the prototype of raccoons, than the skunk prototype it conformed to initially. This shows that typicality, and especially the weighting of typicality, can be determined by something other than, or in addition to, (believed) statistical frequency.

3.3.4 Non-Statistical Typicality

In each of the above cases we find examples of concepts that appear to represent typicalities either in addition to or instead of that which is (at least believed to be) common and distinctive of their referents. This comes out most clearly when we consider what is judged to be their best and most central exemplars. *Type Stability* motivates its claim that concepts represent only that which is typical of their referents on the basis of such judgments: typicality judgments and boundary judgments are what reveal typicality of characteristics. And yet, if this is right, typicality cannot be identified with prototypicality or stereotypicality, the (believed) statistical frequency of characteristics among a concept's (known) referents. Something more, or else, must lie behind the typicalities represent by the concepts posited by *Type Stability*.

3.4 Contextual Determinants of Typicality

In much of the empirical literature on typicality and its hallmarks—typicality effects and boundary effects—there has emerged a persistent theme. Typicality is *context-sensitive*.⁶³ That is, judgments concerning the typicality and boundaries of a concept are subject to the influences of contextual factors: what is judged more or less typical of a concept's referents in one context can diverge from such judgment in another, and which exemplars of a concept are judged to more or less central can diverge across contexts in much the same way. But if this is right, concept typicality cannot be determined entirely by the (believed) stable statistical frequencies of the characteristics of a concept's referents, for such statistical frequencies do not vary across contexts. Contextual factors thus appear to have a role to play in typicality too. And to this extent, the problem is not simply for statistical interpretations of typicality, but also for the idea that concepts represent a stable set of typical characteristics, i.e. the claim of *Type Stability* in general.

3.4.1 Situations

A particularly influential early result is that of (Roth & Shoben, 1983), who showed that sentences that made particular sorts of *situations* especially salient could affect the typicality ratings of a concept. They found, for example, that if participants were primed with a sentence that emphasized a *riding* situation, *horses* and *mules* were judged to be more typical of ANIMAL, but if they were primed with a *milking* situation, *cows* and *goats* were judged to be more typical.

⁶³ See the surveys in (Barsalou L., 1987) and (Dieciuc & Folstein, 2018). And for the context-sensitivity of conceptual cognition in general see (Casasanto & Lupyan, 2015) and (Yee & Thompson-Schill, 2016), as well as the discussion of context in the next Chapter.

Similar results were reported by (Anderson, et al., 1976), who found, for example, that if participants were primed with sentences that emphasized a *culinary* situation, *chickens* and *turkeys* would be judged to be the most typical referents of BIRD.⁶⁴ And (Barclay, Bransford, Franks, McCarrell, & Nitsch, 1974) report that participants would weight most heavily the *heaviness* of the referents of PIANO in situations that emphasize *moving* a piano but weight most heavily their *nice-sounding-ness* in situations that emphasize their *musicality*.

Roth & Shoben especially offer a rather radical interpretation of such results:

[concept typicality] undergoes a complete restructuring once context is introduced [...] the typicality ordering obtained without context no longer plays an important role once context is introduced. (Roth & Shoben, 1983, p. 369)

Subsequent authors have been more cautious. (Hampton, Dubois, & Yeh, 2006), for example, found that typicality ratings could remain quite stable across different sorts of contexts, but also found that particularly familiar situations could have an impact on boundary judgments. They found, for example, that appliances such as *refrigerators* and *dishwashers*, which are usually considered borderline referents of FURNITURE, would more often be excluded from the category in *retail* situations than in other situations, presumably because retail spaces tend to exhibit appliances in an area separate from the *tables*, *chairs*, *sofas* and so on.

⁶⁴ See also (Anderson & Ortony, 1975), (Halff, Ortony, & Anderson, 1976), (Bisanz, LaPorte, Vesonder, & Voss, 1978), (Garnham, 1979), (Potter & Faulconer, 1979), (Anderson & Shifrin, 1980), (Tabossi & Johnson-Laird, 1980), (Barsalou L., 1982), (Tabossi, 1982), (Whitney, McKay, Kellas, & Emerson, 1985), (Greenspan, 1986), and (Hampton, Dubois, & Yeh, 2006).

3.4.2 Perspectives

Another early result, also reported by (Barsalou L. , 1987), is the influence of *perspective*. The same cognizer can take diverging perspectives with respect to the same concept and produce correspondingly divergent typicality judgments. (Barsalou & Sewell, 1984), for example, found that North American participants would judge different birds as most typical of BIRD while taking different nationality perspectives. Taking (their own) North American perspective on birds, they judged *robins* and *eagles* as most typical. But, when taking (what they took to be) a Chinese perspective, they judged *peacocks* and *swans* to be most typical.⁶⁵

In a series of similar results, (Braisby & Franks, 2000) report that participants would vary both their typicality judgements and their boundary judgements on the basis of taking different perspectives.⁶⁶ They found, for example, that participants would vary their boundary judgments with respect to EGG on the basis of taking the perspectives of a sculptor, a biologist, and a caretaker. Participants judged that only a sculptor would allow the inclusion of *Easter eggs* and only a caretaker would allow the inclusion of *scrambled eggs*, while a biologist would include neither. In a similar example, they compare the perspective that can be taken with respect to LION:

[...] classifying a statue of a lion as a "lion" from the perspective of someone who wishes to draw such an animal (and so needs to differentiate lion statues from horse or sphinx statues [...]) is not in conflict with also classifying the same object as a "non-

⁶⁵ It is important to note that Barsalou and Sewell did not survey Chinese participants to confirm that their North American participants were accurate when taking this "Chinese" perspective. Indeed, Barsalou and Sewell found that different demographics differ in their ability to accurately project the perspectives of others. (Barsalou L., 1987), for example, reports that while faculty members, grad students, and undergraduates were highly reliable at taking each other's perspectives, the ability of people to take the perspectives of their close friends appeared rather poor.
⁶⁶ See also (Franks & Braisby, 1990), (Braisby, 1993), (Franks, 1995), (Braisby, Franks, & Hampton, 1996), and (Braisby, Franks, & Harris, 1997).

lion" from the perspective of someone who is hoping to catch a wild beast escaped from the zoo [...] The classifications are simply effected from different perspectives. (Braisby, Franks, & Harris, 1997, p. 172)

As above, Barsalou interprets such findings as providing strong evidence in favor of the contextual instability of concept typicality. However, (Hampton, Dubois, & Yeh, 2006) again tempers this judgment. It is only particularly familiar situations, with clearly associated perspectives, that can have any robust contextual effect on typicality and boundary judgments.⁶⁷

3.4.3 Familiarity

Another early indicator of contextual influence comes from results reported by (Barsalou L., 1985) and (Barsalou L., 1987).⁶⁸ He found that exemplars that are more *familiar* to cognizers tend to also be judged more typical of a concept than its less familiar referents. Familiar exemplars are, in this context, those that a cognizer subjectively considers to be the exemplars they cognize as referents of a concept most frequently.⁶⁹ That is, although *chickens* might be a much more common referent of BIRD, and known to be so, insofar as they are one of the most populous bird species alive, *robins* and *blue jays* might still be much more commonly cognized *as birds* by a particular cognizer, and thus judged to be more typical of the concept. Barsalou reports results that confirm the effect of such familiarity, even when prototypicality is partialed out, in several

⁶⁷ Braisby, Franks, and colleagues interpret such findings as having no direct bearing on the representational content of concepts, although they do argue that it shows the through-and-through context-sensitivity of conceptual cognition—see especially discussion in (Braisby, Franks, & Harris, 1997). Theirs is thus a "concept-external" account of context-sensitivity, which I discuss and dismiss in the next Chapter.

⁶⁸ See also (Ashcraft, 1978a), (Ashcraft, 1978b), (Glass & Meany, 1978), (Hampton & Gardiner, 1984), (Malt & Smith, 1984), and (Janczura & Nelson, 1999).

⁶⁹ This factor is thus also sometimes referred to as "frequency". Other times, "accessibility" or "dominance".
different concepts, including BIRD, FRUIT, VEHICLE, and CAMPING EQUIPMENT. He again offers a rather radical interpretation of such results:

[...] the determination of graded structure is highly flexible phenomenon. There is no single determinant such as similarly to central tendency that is universally responsible for graded structure. (Barsalou L., 1987), p. 101

But in this case, congruent results continue to be reported in more recent literature. (Abbott & Kemp, 2020), for example, having surveyed over 600,000 reported bird sightings in Oaxaca, Mexico, found that familiarity correlated with typicality in five out of six bird categories: VULTURE, PIGEON, HAWK, FINCH, and WREN, but not OWL. Speculating on the divergence in the case of OWL, they remark that the *great horned owl*, which is frequently sighted, is also considered an ill omen in the area, which could explain the deviation in the correlation between familiarity and typicality (if correct, this would indicate another potential contextual factor that may influence typicality judgments). They also note another example of the influence of familiarity on typicality: *nocturnal* birds are generally considered to be less typical of bird than *diurnal* referents, but they are also much less familiar, since they are much less often sighted.

3.4.4 Expertise

Situational influence and familiarity can also have a compounded effect whereby referents that are cognized most frequently in a particular sort of situation can come to dominate typicality and boundary judgments. This what I call the effect of *expertise*: cognizers who are particularly practiced at interacting with and thinking about the referents of a concept in a particular way can develop typicality and boundary judgments for that concept that depart from the judgments offered by less experienced demographics.

Consider again, for example, the results reported in the discussion of *ideal typicalities* above. I noted there that certain tree experts judge tree typicality according to their *height* and *weediness*. But consider now that it is *maintenance workers* and *landscapers* in particular who offered these judgments. Botanical experts, by contrast, had no similar judgment with respect to the weediness of trees. The latter emphasized only the ideality of *height*. But then, it is maintenance workers and landscapers who tend most frequently to weedy trees, not botanists. And the typicality ratings of undergraduates, who had no particular tree expertise, tended in a third direction: predictably, their judgments correlated best with their judgments of *familiarity*.

Similar divergences have been found with respect to BIRD. (Bailenson, Shum, Atran, Medin, & Coley, 2002), for example, found that both North American bird experts and Itza Mayan bird experts judge typicality for BIRD in ways that diverges from the prototypicality of birds, while North American non-experts (undergraduates) offer judgments that align best with bird prototypicality. Moreover, the two expert groups diverged sharply on their judgments: North American experts judged various *small songbirds* as most typical while Mayan experts judged those same birds to be least typical, despite being more or less equally prevalent in both regions. Indeed, Mayan experts reported that small songbirds aren't "true birds", i.e. they are considered to be borderline birds, while their North American counterparts judged the same birds to be most the central exemplars.

The authors discuss two factors that may explain this divergence. First, songbirds are very small and thus hard to spot without binoculars, which are very commonly used by North American birders, but hardly ever used by Mayan experts. They may thus be less familiar in the Mayan

context. Second, is a cultural consideration. For Mayan experts, birds are culturally valuable in relation to *sustenance*—birds are for eating. But for North American experts, the cultural value of birds is much more *aesthetic*—they are for observing. *Songbirds* are nice to look at and to listen to, but they are rather too small to bother with for hunting and eating. And so, the situations in which these different expert groups most frequently engage with birds diverge, resulting not only in divergent typicality ratings and boundary ratings but also in diverging goals and corresponding ideals that correspond to these divergent judgments.⁷⁰

3.4.5 Contrasts

As we've seen, statistical interpretations of typicality hold that typicality is determined by what is both common and distinctive of referents. And distinctiveness is determined by the prevalence of characteristics among non-members. But the idea seems to be something rather more qualified than this bare explication of distinctiveness may suggest. It is not the case that a property is distinctive of a concept's referents insofar as it is frequently absent among *all* non-referents. Rather, distinctiveness often seems to be treated as if it depends only on a property being frequently absent among the referents of closely related *contrast* concepts, typically those that belong to the same superordinate concept at the same level of abstraction. The relevant contrast concepts for BIRD, for example, are other broad animal concepts, like MAMMAL and FISH. And the relevant contrast concepts for OWL are other bird concepts, like RAVEN and ROBIN. And yet, which concepts count as the relevant contrasts might vary by context, and there with the typicality of

⁷⁰ In line with this divergence, Bailenson et al also speculate that the typicality ratings of North American hunters will likely look much more like those of the Mayan experts than those of the North American birders.

their referents. The relevant contrast concepts for FRUIT, for example, might be other food concepts in some contexts (e.g. MEAT and VEGETABLE) but it will be other plant part concepts in other contexts (e.g. ROOT and FLOWER). And so, what is quite distinctive of its referents with respect to one context may be rather indistinctive in another.⁷¹ For example: *edibility* can only be distinctive of FRUIT in non-culinary contexts, since (almost) all non-fruit foods are edible too.

3.4.6 Contextual Typicality

In each of the above cases, we have evidence of contextual factors interacting with judgments of typicality and judgments of referential boundaries. And since statistical frequency does not vary by context in the same way, we have here a second reason to reject the identification of typicality and statistical frequency. Statistical frequencies cannot be the sole determinants of concept typicality. We also have strong reason to reject *Type Stability* in general, insofar as it holds that the typicalities represented by concepts are stable across different contexts.

3.5 Against Type Stability

The primary motivation for *Type Stability*, i.e. the view that concepts represent only that which is stably typical of their referents, is that concepts admit of both typicality effects and boundary effects: participants consistently judge some referents to be better exemplars of a concept

⁷¹ I haven't encountered any empirical evidence to back up these speculations, but on the cognitive role of contrast in various contexts, see for example (Goldstone, 1996), (Ameel & Storms, 2006), (Levering & Kurtz, 2006), (Davis & Love, 2010), and (Palmeri & Nosofsky, 2010).

than others, and participants also consistently judge that some concepts admit of borderline referents. And yet, the primary interpretation of such typicality continues to be a statistical one. Theorists continue to assume that typicality is equivalent to some statistical notion like prototypicality or stereotypicality. I've discussed here two reasons to deny this assumption. First, we have examples of concepts that admit of non-statistical typicalities, like ideal typicalities, normative typicalities, and dependence typicalities. And second, I have presented some evidence for the context-sensitivity of typicality, wherein non-statistical factors can influence both typicality judgments and boundary judgments across different contexts. This evidence places pressure on *Type Stability* in two respects: first, it suggests that the typicalities represented by concepts aren't always *statistical*, and second, it suggests that the typicalities represented by concepts aren't always *statistical*.

Type Stability faces two options. First, it can retain the statistical notion of stable typicality, and explain the bodies of evidence surveyed here by appeal to some *concept-external factors*. This is the route favored by several contemporary proponents, including (Hampton, Dubois, & Yeh, 2006) and (Dieciuc & Folstein, 2018). On their view, concepts represent only stable statistical typicalities, which can be subsequently *augmented* by contextual factors that include not only the contextual factors discussed in §3.4, such as situations, perspectives, contrasts, and the rest, but also the phenomena explicated in §3.3. They treat ideals, norms, and theories of essence and dependence as external contextual factors that can work to influence typicality and boundary judgments in particular contexts but which would otherwise be determined solely by statistical frequencies absent any specific contextual factors.

The second option is to give up the statistical notion of typicality and to adopt instead a *variabilist* notion of typicality according to which that which is typical of a concept's referents can

vary across contexts according to various factors including, but not limited to, various and varying ideals, norms, theories, as well as various and varying situations, perspectives, familiarities, expertise, and so on. Coupling this notion of typicality with the basic idea that concepts represent that which is typical of their referents, we arrive at the position I call *Concept Variability*. According to this view, concepts do represent the typical characteristics of their referents, but what is typical of a concept's referents is can vary, and so, what characteristics concepts represent of their referents can vary. This is to deny then both *Type Stability*, according to which concepts represent stable types, and *Criterial Stability*, according to which concepts represent stable criteria. Indeed, adopting *Concept Variability* constitutes a wholesale denial of *Concept Stability*, according to which concepts represent stable contents of any sort. But, in order to vindicate *Concept Variability*, I must argue against the former "concept-external" approach that may entice the proponent of statistical typicality. Such "concept-external" strategies, along with one other rather radical interpretation of the apparent variations, are my target in the next Chapter.

4.0 Making Sense of Content Variation in Concept Use

4.1 Introduction

As we saw in the previous Chapter, empirical studies into the use of concepts have shown that the representational contents of concepts appear to interact, sometimes rather dramatically, with the contexts of their use, in a variety of forms.⁷² What is considered characteristic of a concept's referents in one context may be judged rather more atypical in others. What is considered an exemplary referent in one scenario may be deemed quite borderline in another. Furthermore, whether some object is judged as belonging among a concept's referents can vary as the context of that concept's use varies, as can the characteristics cognizers are willing to attribute to those referents. And all this variation occurs not only *between* different cognizers in and across different contexts, but also *within* the same cognizers across different contexts of a concept's use. The contents of concepts thus appear to vary both inter-personally and intra-personally across a large variety of concepts implicated and employed in our everyday cognitive acts.

The ever-growing body of data is compatible with several interpretations, two of which have been defended and disputed in the recent literature.⁷³ In line with the view I call *Concept Stability*, (Machery E. , 2015) argues that concepts should be understood as storing a *stable* set of representational contents, which are activated "by default", i.e. "in a context-independent manner" across all the contexts of a concept's use, by all cognizers competent in that concept. On this

⁷² See, for example, surveys of this evidence in (Barsalou L., 1982), (Barsalou L., 1987), (Smith & Samuelson, 1997), (Casasanto & Lupyan, 2015), (Yee & Thompson-Schill, 2016), and (Lohr, 2017). And for the typicality data in particular, see (Dieciuc & Folstein, 2018).

⁷³ Additional data, along with a presentation of the debate is provided by (Lohr, 2017).

interpretation, whatever contents are activated only in some, and apparently in virtue of, contextual factors, and may thus vary across the different contexts of a concept's use are to be understood as contents that are contextually *cued* rather than *stored* within the activated concept. That is, only those contents that are activated regardless of context are to be understood as actually belonging to the representational contents of the activated concept. All other contents, i.e. all contents that are activated only on occasion, under only certain circumstances, only in particular contexts, are to be understood as being *external* to a concept's content. This view has been dubbed *Invariantism*, reflecting its claim that the representational contents that belong to a concept are *invariant* across different cognizers and the different contexts of its use.

According to (Casasanto & Lupyan, 2015), meanwhile, the evidence is to be interpreted in exactly the opposite manner. On their view, concepts represent *only* context-dependent contents. On their view, there is no such thing as default, or context-free, representational contents. All contents are contextually cued. Moreover, all the contextually cued contents of a concept belong to the content of that activated concept. Their view is dubbed, accordingly, *Contextualism*, reflecting the claim that the contents that belong to a concept are *contextually cued*, and thus may vary across different contexts, as well as across different cognizers.

But the difference between these two views goes further. Only *Invariantism* allows that concepts can (and usually do) *persist* across contexts. According to *Contextualism*, by contrast, concepts exist only while and to the extent that they are active. On this view, concepts come into existence in and for a particular context, representing their contextually cued contents, before going back out of existence as soon as they have finished their contextually determined cognitive purpose. As (Casasanto & Lupyan, 2015) put it, "all concepts are *ad hoc* concepts". No concept exists prior to the context that motivates its creation, and no concept survives the end of its

contextually determined use. No concept is re-used, and no concept is shared. In this respect then, *Contextualism* too commits itself to *Concept Stability*, insofar as concepts only every represent one body of contents, contents which trivially cannot vary across contexts, for no concept is used in more than one context.

In opposition to each of the views, I prefer an interpretation of the evidence that denies *Concept Stability*. On my view, the contents of concepts are not stably represented across the various contexts of their use. Instead, the representational contents of concepts can vary. In opposition to *Contextualism*, I thus argue that concepts can (and usually do) *persist* across the different contexts of their use, and thus that they can be re-used by the same cognizer as well as shared between different cognizers.⁷⁴ But, in opposition to *Invariantism*, I hold that concepts need not represent the same contents across the different contexts of their use. Different contexts, according to the particular circumstances of that context. No contents are activated "by default", even if some contents are activated rather often. Furthermore, I hold that the store of contents, from which activated contents are selected, can *change* over time. According to this thesis then, which I call *Concept Variability*, a concept's store of representational contents is *dynamic*, changeable over time, as well as *selective*, activated only sometimes in and according to different contexts of a concept's use.

This intervention on the literature is significant in two respects. First, by proposing *Concept Variability*, I articulate a third option in a debate that has thus far been considered uncontroversially binary. The presiding presumption so far has been that an argument for *Invariantism* is *ipso facto*

⁷⁴ I do not hereby intend to deny the existence of ad hoc concepts. Perhaps we do sometimes devise new concepts for contextually determined use. I just deny that cognition in general is facilitated by such concepts. Ad hoc concepts, if they do exist, exist alongside persisting concepts. For discussion of such concepts see, for example, (Barsalou L., 1983) and (Carston, 2002).

an argument against *Contextualism*; that evidence for *Contextualism* is *ipso facto* evidence against *Invariantism*. And indeed, this diametric division makes sense if we hold fixed the commitment to *Concept Stability*. For if concepts represent a stable set of contents, invariant across different contexts and cognizers, then concepts can be individuated by their contents. And that means that evidence for the *persistence* of a concept is *ipso facto* evidence for the *invariance* of its content, while evidence for the *variability* of content is *ipso facto* evidence for the *ephemerality* of concepts. Denying *Concept Stability* opens up a third option, encapsulated by *Concept Variability*: persisting concepts representing variable contents.

But second, and most importantly, the third option offered by *Concept Variability* offers a more satisfactory interpretation of the evidence than each of the two existing options. That is, *Concept Variability* offers not only a new option, it offers a better one, one which makes better sense of the both the evidence in favor of variable content and the evidence in favor of persistent concepts.

I proceed as follows: I start in §4.2 with some preliminaries, articulating some of the assumptions that I share with both *Contextualism* and *Invariantism*, assumptions that have been left at least somewhat implicit in the preceding Chapters. These include the idea that concepts are to be understood as *explanatory posits*, that they are posited to explain our *cognitive behaviors*, and that concepts represent both *epistemic contents* (bodies of general knowledge) and *intentional contents* (categories of referents), which together constitute their representational contents. In §4.3 I review some of the evidence provided by our cognitive behaviors that constitute the *explananda* that I argue *Concept Variability* explains better than both *Invariantism* and *Contextualism*. This includes, but exceeds, the evidence I considered in the previous Chapters. In §4.4 I present and reject the *Invariantist* alternative, arguing that the explanation on offer succeeds in explaining

concept persistence only by failing to adequately account for the full extent of content variation. In §4.5 I present and reject *Contextualism*, arguing that the explanation on offer succeeds to explain content variation only by failing to adequately account for the evidence for concept persistence. I close in §4.6 by rearticulating the happy middle offered by *Concept Variability*, showing how the denial of *Concept Stability* allows us to make sense of both the persistence and the variation evident in our concepts.

4.2 Preliminaries

One of the standard shared assumptions in the study of concepts is that they *enable cognition*. They are routinely described as the "building blocks of thoughts" (Margolis & Laurence, 2019), implicated in an enormous variety of cognitive processes, including "memory, learning, and decision-making" (Margolis & Laurence, 2019), communication (Murphy G., 2002), and "most, if not all, higher cognitive competences" (Machery E. , 2009, p. 4). As (Marques & Wikforss, 2020, p. 1) put it, "Concepts stand at the center of human cognition. We use concepts in categorizing objects and events, in reasoning and action, and in social interaction." Indeed: "Without concepts, there would be no thoughts" (Prinz, 2002, p. 1).

There are several ways of unpacking what this enabling relation might involve. I understand it as an *explanatory relation*: we posit concepts in order to *explain* our cognitive behaviors. It is *because* we possess a concept of some category, say the concept FRUIT, that we are able to engage in various cognitive processes that concern fruit. It is competence in this concept that allows us to think and talk about fruit, identify and classify fruit, perform both theoretical and

practical inferences that concern fruit, reliably assent to and demur from different utterances that involve reference to fruit, and so on.

According to this assumption then, the ability to perform particular cognitive acts provides evidence for the possession of some particular concept. But these abilities also provide evidence for the *contents* of our concepts. And in particular, as I see it, our cognitive behaviors offer evidence for two sorts of representational contents belonging to our concepts.

The first is a concept's characterizing content, or what I'll call its *epistemic content*, which is constituted by the bodies of general knowledge stored by a concept.⁷⁵ If someone tends to think and talk about fruit as if they are *edible* say, asserting and assenting to sentences that predicate edibility of fruit, and perhaps even occasionally *eating* fruit in their environment, buying and chopping up fruit to put in their fruit salads, and so on, we are well justified to infer that their concept of fruit includes the knowledge that fruit are *edible*. This characteristic of fruit is thus part of their concept's epistemic content.

Second is a concept's referential content, or what I'll call its *intentional content*, which is constituted by the categories of objects to which a concept refers. If a cognizer reliably thinks and talks about the category fruit as including *apples*, for example, identifying and classifying apples as fruit, asserting or assenting to sentences such as "Apples are fruit" or "Fruits include apples", sometimes including apples in their fruit salads, and so on, we are well positioned to infer that their concept of fruit extends to at least *apples*. This subcategory of fruit is thus part of the concept's intentional contents.

⁷⁵ I say "epistemic" in order to allow the inclusion of belief-like contents, i.e. knowledge-*that* the referents of a concept are generally like this, that, and the other, as well as ability-like contents, i.e. knowledge-*how* the referents of a concept can generally be employed, interacted with, and so on. And I say "general" in order to allow that the represented knowledge may be or less *exception-tolerant* (i.e. generic), and thus may be true of only some of a concept's referents.

To this extent I am in agreement with both *Invariantism* and *Contextualism*, each of which takes on board the idea that our concepts and their contents are revealed by our cognitive behaviors, and that these behaviors that are to be explained by positing the possession of particular concepts with particular contents, which may include reference to both the bodies of general knowledge that characterize the referents of a concept as well as the categories of those referents themselves.

I depart from the two theses by denying the further assumption of *Concept Stability*, according to which the epistemic and intentional contents of concepts are *stable* across the different cognitive acts that they enable. And my reason for this departure is ultimately rather simple: the evidence provided by our cognitive behaviors across different acts, different contexts, different cognizers, and different disciplinary domains, reveals significant *variation* in both the epistemic and intentional contents of our concepts. While *Concept Stability* predicts that our concepts represent the same objects and the same characteristics across the various contexts of their use, the evidence actually provided by our cognitions indicates that these contents can vary. *Concept Variability* takes this evidence at face value, claiming that the representational contents of our concepts and the sinterpretation, we must first consider some of that evidence.

4.3 Content Variation in Concept Use

In the previous Chapter I discussed some examples of how contexts can interact with typicality judgments, evidence that suggests that cognizers will judge different characteristics of a concept's referents as more or less typical of that concept across different contexts of the concept's use. The same was true of boundary judgments, which offers evidence that cognizers will judge

different objects as more or less borderline of that concept across different contexts. These differences already indicate the possibility of variability in both the epistemic and the intentional contents of our concepts. But there are several additional bodies of evidence, pointing to both *intra-personal* and *inter-personal differences*, that I think worthwhile considering in adjudicating the case against *Concept Stability*, and in favor of *Concept Variability* in its stead. First, I consider some evidence that supports the idea that concepts are *dynamic*, i.e. that the contents that they store can *change* over time. Then I return to some evidence for the idea that our concepts are *selective*, representing only *some* of their store contents on any particular occasion of their use.

4.3.1 Conceptual Development

The early empirical studies into the conceptual contents revealed by our cognitive behaviors focused on adult participants in several more or less artificial experimental contexts. But cognitive scientists soon broadened their view. One early expansion was into the empirical study of *children's* concepts and how they both overlap with and depart from the corresponding concepts of their adult caretakers. The findings were as interesting as they were robust. It has repeatedly been found that children's concepts differed from the corresponding adult concepts in both their epistemic and intentional contents, suggesting not only that these contents may differ between different people, i.e. between children and adults, but also that they can differ within the same individual over time, i.e. over the course of child growing into an adult.⁷⁶

Consider, for example, some findings with respect to FRUIT. In this case it has been found that the *epistemic contents* of children's concepts are much *thinner*, including comparatively little

⁷⁶ See, for example, discussions in (Anglin, 1977), (Carey, 2009), and (Amin & Levrini, 2018).

knowledge as compared to the parents, guardians, and teachers from which these concepts were learnt. For example, (Anglin, 1977) found that children (in the industrialized West, at least) focus very heavily on the *edibility* of fruit, having little knowledge or interest in the fact that fruit *come from plants*. Here, for example, is 4-year-old Sharon ("S") reporting her knowledge of fruit to an examiner ("E"):

E: OK, do you know what fruit is?

S: Yeah.

E: What's fruit?

- S: There's um grapes, bananas, apples, and pineapples, um pears, and let's see, I don't know any more.
- E: OK, and what kind of thing is a fruit?
- S: Um, it's what you eat.
- E: Uh huh.
- S: And it gives juice.
- E: Uh huh.

S: And, and it, makes pies, and it makes some pies that are called apple pies sometimes.

E: Uh huh.

S: And, and we had a, we had a movie about um apples, and, and, it was about Johnny Appleseed. (Anglin, 1977, p. 210)

Notice how she mentions only *food*-related knowledge of fruit: they are edible, and they can be made into juice and pies. Notice also that she starts by listing a few exemplars of fruit, but this list is rather small.⁷⁷ She reports knowing only five types of fruit. Thus the *intentional* contents of her concept of fruit appears to be similarly narrow, at least compared to what her caretakers

⁷⁷ And finally, note that after she mentions apple pies, she starts shifting attention to some of her knowledge concerning *apples* in particular, rather than fruit in general. I read this shift as an indirect suggestion that Sharon thinks of *apples* as some of the most exemplary referents of FRUIT, a tendency she appears to share with adults, as confirmed by both (Anglin, 1977, p. 53) and (Van Overschelde, Rawson, & Dunlosky, 2004, p. 302).

would (presumably) include among the referents of FRUIT. And indeed, Anglin reports that while the adults that participated in his study on average underextended the concept in only 5% of cases, the children underextended the concept in an average of 37.1% of cases. (Anglin, 1977, pp. 111, 117) That is, upon being asked to classify some given object as either a fruit or a non-fruit, adults classified a given fruit as a non-fruit only 5% of the time, while the children classified some given fruit as a non-fruit of the time. Their concept thus included not only *less knowledge*, but it also extended to *fewer objects* than that of adults, showing that both the epistemic and intentional contents of a child's concept FRUIT excludes contents included in the adult's FRUIT.

It is worth noting that children's concepts don't always *underextend* in comparison to the concepts of adults. Although Anglin's findings tend to show more underextension than overextension, one notable exception to this tendency is the concept FLOWER, which Anglin's child participants overextended in an average of 18.8% of cases in comparison to the adult rate of a mere 2.5%. (Anglin, 1977, pp. 111, 116). Children appeared especially prone to applying the concept to several sorts of flowerless plants, suggesting that the younger participants did not have a clear sense that flowers are merely a particular sort of plant-*part*. They certainly show no evidence of knowing that flowers are the sorts of things that eventually give way to fruit. This knowledge will come later.

4.3.2 Conceptual Sophistication

It can be tempting to assume that our initial childhood concepts eventually give way to adult concepts that tend to stay fixed after an initial burst of childhood learning. As Wilson puts the idea: This is the tenet that I call *semantic finality*, <u>viz.</u>, the claim that, with respect to a wide range of basic vocabulary, competent speakers acquire a *complete conceptual mastery* or *grasp* of their word's semantic contents by an early age—no later than 10 or 11, say. This core content then acts as an *invariant* that underwrites many of our characteristic endeavors [...] To be sure, it is conceded that, beyond their initial period of conceptual inoculation, speakers will often tinker with these early basic contents in minor ways—e.g., later we learn that the usage of "dog" can permissibly extend to cover the wider family <u>Canidae</u> and poetically stretched to embrace human feet. Nonetheless, the majority of matters we subsequently learn about dogs—that Jones' specimen down the street is an ugly brute; that they are largely color blind; that they are available in sizes smaller than squirrels, etc.—do not alter the stored core content of *being a dog* and can be ignored by the student of semantics proper. (Wilson, 2006, p. 19)

Against this tenet of "semantic finality", Wilson presents a series of cases, drawn especially from the applied sciences, that show that conceptual changes and sophistications of various sorts can continue throughout adult life, as adults gain expertise in particular disciplines or domains of study, like engineering and mechanics, which deal in such concepts as HARDNESS and WEIGHT. In these cases, Wilson reports a "patchwork" of contents that result from advanced expertise, with disciplinary experts displaying significant nuance with respect to their ability to determine the hardness, weight, and so on of different materials. His account of the complexities HARDNESS is particularly illustrative:

[...] our usage of the predicate "is hard" displays a fine-grained structure that we are unlikely to have noticed, for our everyday usage is built from local patches of evaluation subtly strung together by natural links of prolongation. More specifically, in everyday contexts we adjudicate the "hardnesses" of various materials, both comparatively and absolutely, through a wide variety of comparatively easy to apply tests—we might *squeeze* the material or *indent* it with a hammer; attempt to *scratch* it or *rap* upon it;

and so on. [...] In fact, our choice of tests is likely to have been suggested by the material in question: we instinctively appraise a wood by rapping upon it, a rubber by squeezing, a metal by attempting to make a small imprint; a glass or ceramic by rapping lightly or scratching (not by trying to make a small imprint!). [...] These hidden forms of preferential technique for assessing "hardness" become quite salient as everyday informal methods become improved into the carefully calibrated forms of testing apparatus that go by such titles as Brinell or Vickers indenters (vigorous squeezing and then releasing); superficial Rockwell testing (mild squeezing and partial releasing), durometer (squeezing without releasing), sclerometer (scratching), scleroscope (a different instrument that raps its specimen), the Charpy impact test (hitting with a hammer) and so forth. (Wilson, 2006, pp. 336-337)

Wilson does not present any cognitive data to suggest that such advanced techniques of discerning hardness make any difference to the conceptual contents represented by an expert's concept of HARDNESS and so on. But some congruent findings have been reported with respect to the neuroplasticity in disciplinary experts in other domains. (Beilock, Lyons, & Mattarella-Micke, 2008), for example, report the neural effects of expertise in ice hockey on language comprehension, showing that hockey fans and players have a much more nuanced semantic representations of action than novices. Similar results are reported by (Hoenig, et al., 2011) with respect to musical expertise and the representation of sound. Importantly, these findings show that expertise in motor skills can show up in semantic representations, which thus end up significantly richer than the corresponding representations of novices.

Put in terms of concepts, such evidence suggests that we can keep learning and adjusting our concepts, like that of ACTION and SOUND, updating the contents of our concepts as we continue to move through the world, regardless of age. The same may be said of HARDNESS, WEIGHT, and the like. I read this evidence as indicating the possibility of conceptual changes *throughout* life, which carries the implication that the contents of our concepts can, and often will, differ *both inter-personally*—between children and adults, experts and novices—but also *intrapersonally*—as individuals progress from childhood to adulthood, and as novices advance to greater levels of disciplinary expertise. It also suggests, in line with *Concept Variability*, that our concepts *persist* through a variety of changes in their epistemic and intentional contents.

4.3.3 Inter-Disciplinary Differences

The difference disciplinary expertise can make becomes especially evident once we consider the *divergent* sophistications that can occur across the different disciplines a concept is employed. As Fee and colleagues put it with respect to HARDNESS:

The definition of hardness varies depending upon the experience or background of the person conducting the test or interpreting the data. To the metallurgist, hardness is the resistance to indentation; to the design engineer, a measure of flow stress; to the lubrication engineer, the resistance to wear; to the mineralogist, the resistance to scratching; and to the machinist, the resistance to cutting. (Fee, Segabache, & Tobolski, 1985, p. 71)

I already reported divergences in the typicality concepts of BIRD across different sorts of bird experts and TREE across different sorts of tree experts in the previous Chapter: what North American birdwatchers find most exemplary and typically characteristic of the referents of BIRD depart significantly from what Mayan hunters find most exemplary and typically characteristic, and what landscapers, maintenance workers, and botanists judge most exemplary and typically characteristic of the referents of TREE differ drastically too. Such variations in typicality and boundary judgments reveal the sorts of differences different disciplines of expertise can make. Even more radical divergences can be observed with respect to such concepts as FRUIT. In this case, there are some significant epistemic and intentional differences that can be observed between the use of the concept in *botanical* contexts and its use in *culinary* contexts. As is well known, from the botanical perspective, fruits are understood as the *seed-bearing structures of flowering plants that develop from plant ovaries after flowering*. From this botanical perspective, FRUIT includes amongst its referents a wide variety of seed-bearing structures, including several sorts of edible "vegetables" (e.g. *pumpkins, olives, beans, peas, corn,* and more), certain sorts of edible nuts (e.g. *chestnuts, hazelnuts, peanuts,* etc.), as well as some inedibles (e.g. *acorns, thorn apples, winter berries,* and so on). This is, of course, a much wider conception than is typical of a culinary perspective, according to which fruits are *juicy, crunchy, edible produce,* including such fruits as were listed by little Sharon, i.e. *grapes, bananas, apples, pineapples,* and *pears,* and the like, and typically excluding the above nuts and culinary vegetables.

To the extent that these divergences can be attributed to one concept shared across different individuals at different stages of their intellectual development, across different disciplines of expertise, the divergent sophistications in such concepts as HARDNESS, BIRD, TREE, and FRUIT, further bolsters the idea that our concepts are *dynamic*, altering and updating their contents as we learn and grow, gaining expertise in the many different characteristics or their various referents. It also bolsters the idea that our concepts are *contextually selective*, insofar as the same person might employ the same concept while moving back and forth between different sorts of contexts: consider, for example, a fruit botanist perusing the produce section for some ingredients for a fruit salad.

4.3.4 Intra-Disciplinary Differences

It can be quite tempting to conclude that the above divergences offer evidence that different disciplines are employing *distinct* concepts, despite the continuing overlap in contents and shared vocabularies. I will address this idea in fuller detail below, but for now I think it is worth noting the extent to which divergences appear also *within* domains. That is, there is evidence that substantive differences and disagreements exist over the content of ostensibly shared concepts not only *across* such domains as botany, cuisine, and the rest, but *within* these domains too.

Philosophers of science have long noticed the lack of conceptual monism within different branches of science.⁷⁸ SPECIES is the paradigm case. There are at least three prominent and legitimate non-equivalent conceptions of species of use in contemporary biology, which differ in both their epistemic contents and their intentional contents: the *phenetic* conception of species, the members of which are morphologically similar, the *phylogenetic* conception of species, the members of which share a common natural history, and the *biological* conception of species, the members of which can interbreed. The third conception is easily applied in pursuit of many different research purposes, but it cannot be employed in the case of *asexual* organisms, nor in the case of *hybrids*, and thus has limited scope. The second conception requires more sophisticated research tools to employ successfully and can thus be overly demanding in many research contexts, but it has the advantage of overcoming some of the shortcomings of the biological conception. And yet, it cannot be employed by paleontologists, who have no access to their subjects' genetic

⁷⁸ As (Taylor & Vickers, 2017) make especially vivid, examples include SPECIES (Kitcher, 1984); (Ereshefsky & Reydon, 2015), CONSCIOUSNESS (Block, 1995), (Irvine, 2013), ACID (Hendry, 2008), MEMORY (Baddeley, Eysenck, & Anderson, 2009), GENE (Brigandt, 2010), CONCEPT (Machery E., 2009), and MUSIC (Currie & Killin, 2016), among many others.

material, nor to any direct evidence of their breeding behaviors. In this case the first conception is particularly useful. But of course, shared morphology can be misleading. And so no conception is perfect for all biological purposes, and yet each has its particular strengths.

The same is true of FRUIT. It is clear that some conceptions of fruit are better suited for culinary purposes while others are better suited for botanical purposes. But even in botany alone, there are at least two important and legitimate non-equivalent botanical conceptions of fruit.⁷⁹ One conception *includes* reference to plants parts known as *accessory fruits*, which are ripened plant ovaries that have incorporated adjacent plant matter in the course of their maturation, and are thus only partially constituted by plant ovaries. These include what are typically thought to be the most exemplary of fruits: *apples, pears, figs, pineapples*, and *strawberries*. The second conception, the referents of which are often referred to as the "true fruits", *excludes* such plant parts, extending only to those fruits that are constituted by ripened plant ovaries alone, e.g. *grapes, bananas*, and *pomegranates*. Despite such well-known divergences, both conceptions of fruit remain of use in botany, each fulfilling some research purposes, with neither fulfilling all.

If we are to read the evidence for *inter*-disciplinary differences as indicating the existence of distinct concepts across different domains, then the evidence for such *intra*-disciplinary differences, should be read as offering the same indication: there are multiple distinct concepts employed *within* the sciences for such subject matters as hardness, species, and fruit. Where we might have expected *one* concept, we find instead *several*, and indeed sometimes we might find *very many*. Better, I think, to understand these as very many *variants* of one shared concept, which can be shared across different life-stages, different levels of expertise, different more or less interrelated research programs within a single discipline, and different more or less inter-related

⁷⁹ See, for example, the explication and discussion of the difference in (Esau, 1999).

disciplines and domains. Indeed, I think we can say as much for the divergences and disagreements evident across different more or less interrelated and interacting cultures.

4.3.5 Cross-Cultural Differences

Since so many of our empirical studies into the contents of many different concepts draw primarily from the observed cognitive behaviors of Western industrialized peoples, and especially from the observed cognitive behaviors of North Americans, it can sometimes be obscured just how much variation and difference there exists between how people in different cultures conceptualize the same subject matters. Nevertheless, significant cross-cultural differences are quite easy to see once you start looking.

Consider again the case of FRUIT, the culinary usage of which can differ across different culinary cultures and cuisines. This can easily be observed in the nutritional guides developed by nutritionists and cultural experts in different nations for local populations.⁸⁰ These guides reveal significant regional differences in exactly which sorts of botanical fruits are included or excluded from the recommended culinary fruit.⁸¹ Some guides, for example, explicitly exclude *legumes* from their recommended fruits (e.g. South Africa, Bangladesh) while others include them (e.g. Benin). Some exclude *nuts* (e.g. Sierra Leone, Afghanistan), while others include them (e.g. Seychelles, Germany). Most also include at least some botanical fruits among their recommended vegetables, although which these are varies regionally too (e.g. Nigeria includes *avocados* among

⁸⁰ All guides can be accessed at <u>https://www.fao.org/nutrition/education/food-dietary-guidelines/regions/en/</u>. (Accessed Oct 27, 2021.)

⁸¹ Indeed, some nations produce multiple guides in order to best reach different demographics among their populace. Australia, for example, provides a separate guide for indigenous Australians, recommending more indigenous foods and fewer imported Western foods, including fruits, more familiar to settler communities.

the fruits, but Indonesia puts them with the vegetables, while Cuba places them with the fats). Other guidelines simply group all edible produce together, marking no distinction between culinary fruit and vegetables (e.g. Namibia, Canada, China), but of course which fruits and vegetables these guides recommend vary by region.

Such cross-cultural differences are evident in other concepts too. As I'll detail in the next Chapter, the many varied ways different communities, disciplines, and institutions across the world think about fish reveal remarkable variations, in different fish-studying corners of ichthyology, in different fish-eating culinary cultures, in different fish-harvesting and fish-trading jurisdictions, and more.

Further interesting differences have also been reported for various basic concepts such as SPACE (Majid, Bowerman, Kita, Haun, & Levinson, 2004), TIME (Boroditsky, 2001), MOTION (Papafragou, Hulbert, & Trueswell, 2008), COLOR (Regier & Kay, 2004), FLAVOR (Osawa & Ellen, 2014), BODY PART (Enfield, Majid, & van Staten, 2006), and different body part concepts, like that of ARM (Brown, 2013), among many others.

In some cases the reported differences are in the concept's *intentional* contents. For example, (Brown, 2013) reports that in about a third of over 600 linguistic groups surveyed, arms are thought to *include hands*, and thus to extend from shoulders to fingertips, while two thirds think arms *exclude hands*, and thus end at the wrist. And some populations mark no sharp cutting off point between hands and arms, with hands gradually ending and arms gradually starting as you move up the forearm, with the middle forearm existing as a borderline arm/hand. These different linguistic groups thus appear to think of different parts of the human body as included among the referents of ARM.

Other differences are in the *epistemic* contents of concepts. (Boroditsky, 2001), for instance, found that while cognizers fluent in English tend to think of time as moving *horizontally*, left to right, cognizers fluent in Mandarin tend to think of time as moving *vertically*, top to bottom, differences that clearly emerge in how the concept is employed in experimental categorization and inference tasks. For example, English-speaking participants were significantly slower in categorizing times when prompted to conceptualize them as progressing vertically, while Mandarin-speaking participants displayed the opposite tendency.

Once again, we might be tempted to read these differences as indicating that different cultures are employing different concepts for the same subject matter. This would, of course, be in addition to the different concepts employed by different members of these cultures: the children, the novices, the many different experts, in and across different domains of expertise, at different periods of their development, education, and training, all employing distinct concepts of fruit, fish, arms, time, flower, hardness, species, birds, trees, and more. I think this is a mistake. But before I defend my reading of the evidence, we must return to one more body of evidence, that which indicates the intra-personal differences that can emerge once we look at how a concept is employed by a single person across different occasions of its use.

4.3.6 Cross-Contextual Differences

So far, I've been offering evidence in favor of the idea that the contents of our concepts are *dynamic*, and thus can persist through change. Both the bodies of knowledge and the categories of objects represented by some given concept can change over time, as we learn, grow, and move through different domains and regions of the world. This is one sort of content variation: variation in which contents our concepts *store* over time. But I've also offered evidence for the idea that

these changes can develop in different directions on the basis of different levels and types of disciplinary expertise, as well as different levels and kinds of familiarity with different cultural practices and norms. To the extent that a single person may have access to a plurality of these forms of expertise, such differences already suggest the possibility of another form of variation, i.e. the variation in *which* of a concept's stored contents are *selected* for use in cognition on any given occasion of a concept's use. Time to consider such selective representation directly.

Evidence for contextual selection of contents emerged quite early in the empirical study of conceptual contents. (Anderson, et al., 1976), for example, found that their participants were more likely to list *grapes* as exemplars for FRUIT if cued by a sentence like "The fruit was made into wine" than by "The fruit was served with the wine." Similarly, (Greenspan, 1986) found that participants were more likely to associate *climbability* with TREE if cued by a sentence like "The children played in the tree" than by "Henry chopped up the tree." Recall also the results reported with respect to typicality judgments in the previous Chapter: (Roth & Shoben, 1983) found, for example, that if participants were primed with "Fran pleaded with her father to let her ride the animal", *horses* and *mules* were judged to be more typical of ANIMAL, but if they were primed with "Stacy volunteered to milk the animal whenever she visited the farm", *cows* and *goats* were judged to be more typical.

Further intra-personal instabilities were found in studies that surveyed participants across different days. Consider, for example, Bellazza's findings with respect to several concepts, including ANIMAL, FLOWER, FRUIT, INSECT, MUSICAL INSTRUMENT, SPORT, and TREE. When asked to list exemplars of such categories across two sessions spaced one week apart, (Bellezza, 1984a) found that intra-personal reliability averaged about 50%. That is, participants would, on average, list only about half of the referents they listed a week earlier. And, when the

task was to list characteristics attributable to such exemplars, also across two sessions spaced one week apart, (Bellazza, 1984b) found that intra-personal reliability averaged 69%, listing on average just over two-thirds of the same characteristics as they listed in the previous session. There is thus clear *overlap* in the contents being activated across the two sessions, but there are also notable differences, differences that emerge even without any explicit contextual cuing.

Such intra-personal differences don't only concern which contents happen to come to mind in any particular session. We also contradict our earlier judgments. (McCloskey & Glucksberg, 1978), for example, found that their participants would regularly change their classification decisions across sessions, especially with respect to less typical referents.⁸² With respect to FRUIT, for instance, they found that roughly a third of their participants changed their classification decisions for less typical referents: 27% changed their decision for *pumpkins* and 33% changed their decision for *olives*. Once again, such differences emerged even without explicit contextual cuing.

Further evidence concerns the classification of different types of referents across different contexts. Different standards can be operative for the classification of different referents, or for the same referents across different contexts, suggesting that which of a concept's contents drive categorization on any particular occasion can vary across different circumstances. Examples here include RED (Halff, Ortony, & Anderson, 1976), LIE (Coleman & Kay, 1981), MOTHER (Lakoff, 1987), WATER (Malt, 1994), and DESIRE (Jerzak, 2019).

As (Lakoff, 1987) details, whether or not a person counts as a referent of MOTHER, for example, can depend in some cases on whether they are a child's *legal guardian*, in other cases on

⁸² They found that decisions about highly typical and highly atypical referents were comparatively more stable, though by no means *invariant*.

whether they are a *primary caretaker*, and only sometimes on whether they *birthed the child*, or *contributed their genetic material*, or *married another parent*, etc. Similarly, as Malt showcases, whether a particular liquid counts as water can require different levels of H₂O content for different types of liquids: *tea* tends to be judged to have a higher H₂O content than *seawater*, and yet only the latter tends to be judged a referent of WATER.

Such bodies of evidence, which have been replicated and expanded on for a variety of concepts in the decades since, strongly suggest the idea that the contents of our concepts are contextually unstable, apt for selective representation across the different contexts of a concept's use.⁸³ We do not, on every occasion of some concept's use, activate for representation the same objects and the same characterizations. Instead, we can represent different parts of concept's content, and which contents these are can be cued more or less explicitly by the contexts in which the concepts are used.

As I see it then, the available evidence concerning concept use shows that our concepts represent variable contents in at least two respects: not only do the contents stored by our concepts change over time, the contents that are selected from storage for cognitive use on any given occasion can also vary across the different contexts of a concept's use. That is, there are no contents that are stably represented by a concept across all the many times, places, people, and purposes in, by, and for which it is employed. If it really is true that the contents of our concepts are revealed by our cognitive behaviors, and the large body of evidence provided by our cognitive behavior is

⁸³ Again, see the surveys and commentaries offered in (Barsalou L., 1982), (Barsalou L., 1987), (Smith & Samuelson, 1997), (Casasanto & Lupyan, 2015), (Yee & Thompson-Schill, 2016), and (Lohr, 2017). And for the typicality data in particular, see (Dieciuc & Folstein, 2018).

that the contents represented by some given concept across different contexts can vary, then it follows that the contents of our concepts can vary. The contents of our concepts are not stable.

4.4 Against Invariantism

4.4.1 Invariantism

The idea that our concepts represent at least some of their contents stably is more or less standard across large parts of philosophy. And this endorsement of *Concept Stability* includes some philosophers well-acquainted with the empirical evidence that concepts appear to allow remarkable inter-personal and intra-personal variations in their contents. Such views, which have come to be known as *Invariantism*, acknowledge the need to make sense of apparently variable concepts, but decline to admit that the contents of concepts can actually vary. On their view, the contents of concepts really are *invariant*. The variation observed in empirical studies of cognitive behavior is to be explained by something other than an appeal to the contents of our concepts; they are to be explained by some *concept-external* factors.

The basic move in explaining content variation while adhering to such conceptual invariance is to posit a distinction between two sorts of content, only one of which is represented by a concept *itself* across the various occasions of its use. These are (i) *invariant contents*, which are invariantly represented by a concept across its various uses, and (ii) *auxiliary contents*, which can vary, but which are merely associated with a concept, represented only in some circumstances

of its use.⁸⁴ And according to *Invariantism*, a concept's content is exhausted by its invariant content; whatever content variation there may be in found in our cognitive behaviors is due entirely to variation in auxiliary contents.

One of the great advantages of adopting *Invariantism* is that it allows for the idea that concepts can *persist* and that they can be *shared*. For, as long as the invariant contents of some concept endures through various updates, adjustments, and changes in its associated auxiliary contents, the concept persists, and as long as the same invariant contents are shared between different people, across different stages of their development, education, training, and so on, no matter how different and diverging their associated auxiliary contents are, the concept is shared.

The details of such an account can be spelt out in different ways. Medin and Ortony, for example, contend that concepts provide *stable inputs* to cognitive processing, which may produce varying cognitive outputs on the basis of variations in auxiliary inputs. Objecting to my way of reading the evidence, they contend:

[...] we think care has to be taken not to equate instability in outputs or behaviors with underlying or internal instability. Might it be that underlying concepts are in fact stable (whatever that might mean) and that the apparent instability is an artifact of the processes that operate on these stable representations? (Medin & Ortony, 1989, p. 191)⁸⁵

⁸⁴ These auxiliary contents have gone by several different names in the literature: "context-dependent properties" (Barsalou L., 1982), "non-essential information" (Rey, 1985), "heuristics" (Keil F., 1994), "background knowledge" (Machery E., 2015), and "associative characterizations" (Camp, 2015), among others.

⁸⁵ Displaying a similar sort of caution, (Keil F., 1994, p. 170) warns, "Categorization, however, must not be equated with heuristics and other procedures that provide rough and ready identification of instances above a modest confidence level. Although one might use hair length as a rough means for identifying human sex at a certain confidence level, a careful and deliberative categorization of humans by sex would make little note of such an attribute. Careful, considered judgements of membership may emphasize different aspects of mental structures than the fastest and loosest means of identifying members of categories." I, of course, think that *both* these sorts of identification procedure ought to be considered as indicating a part of a concept's content, to the extent that that concept and its contents is what serves to enable the classificatory cognitive act.

They propose that one such auxiliary input is *context*, such that invariant concepts coupled with varying contextual cues produce the representation of various auxiliary contents, which result in varied cognitive behaviors. Thus, the phenomena of content variation can be explained without an appeal to variation in a concept's own content.

The most prominent recent elaboration of such an idea, defended by Machery, understands invariant content as *automatic content*, which is content that is activated across all contexts of a concept's use both quickly and regardless of contextual cues, while context cues the addition of different auxiliary contents on different occasions according to different contexts, where such *contextually cued contents* are cued for use comparatively slowly and only occasionally. As he puts it:

Invariantism holds that in all contexts, automatically, and quickly, we retrieve a specific body of knowledge—say, the knowledge that dogs are four-legged furry animals that bark, are good companions for people, and have a specific shape—and this body of knowledge *constitutes the concept* of dogs. Everything else we know about dogs (e.g., that dogs are mammals and vertebrates, that dogs and wolves are related, that I have a dog called *Fido*, that Labradors are dogs, etc.) is part of the background knowledge about dogs: this knowledge is retrieved in a context-dependent manner. (Machery E., 2015, p. 571)

That is, according to *Invariantism*, some contents are *privileged* for inclusion in a concept. These privileged contents are those that are represented regardless of context. All other contents are *renegaded*, excluded from a concept's contents. These renegaded contents are those represented only in some contexts.

Now, it is worth noting from the outset that which contents are automatically represented can differ between different people as well as the same people over time. Machery admits that what content is represented automatically can be *contingent* and *idiosyncratic*. Not everyone will automatically represent dogs as being good companions for people, for example. And indeed, some automatically represented contents might be incompatible with one another. Some people might automatically represent dogs as *dangerous*, for example. In such cases, Machery must claim that what might seem to be the use of a single concept is rather to be understood as the use of *distinct* concepts, each representing only one of each of the incompatible automatic contents.⁸⁶

For Machery, *Invariantism* thus predicts that content variation exists wherever there exists either some difference in contextually cued content along with invariance in automatic content, or some difference in automatic content belonging to distinct concepts. More generally, *Invariantism* is committed to explaining content variation as follows: wherever we find content variation, it is due to either differences in auxiliary contents or distinctness of concepts. No variation is to be explained by an appeal to the contents of concepts themselves.

4.4.2 The Problem(s) with Invariantism

As I see it, there are at least two problems facing the *Invariantist* interpretation, each of which concerns the account's ability to provide an adequate explanation for the range of phenomena detailed in the previous Section. The first is, as usual, my complaint that the view renders concepts objectionably inert with respect to explaining our varied cognitive behaviors. The second is my concern that *Invariantism* risks empirical invalidity, insofar as its only available

⁸⁶ One can also infer this claim on Machery's behalf on the basis of the views espoused in (Machery & Seppälä, 2010). I will discuss the (im)plausibility of this move in the next sub-section.

strategy for denying variable contents in concepts renders the view unfalsifiable. I take these issues in turn.

4.4.2.1 The Problem of Limited Cognitive Import Again, Again

My first worry for *Invariantism* emerges from the distinction offered between automatic and contextually cued content. The idea that some of a concept's contents are represented regardless of context certainly makes sense in the abstract, but it's hard to see how to employ the notion in practice. How are we to diagnose for some given content whether it is automatic, and thus apt for inclusion in a concept's content? How are we to confirm, for example, Machery's suggestion that being a *good companion* is automatic for the use of DOG? How might we determine whether contents such as *produce* or *seeded* are automatic for the use of FRUIT?

Our best bet is to employ the idea that invariant contents are represented across, as (Machery E., 2015, p. 571) says, "all contexts" of a concept's use. They are thus those contents that are both *stored* in a concept across all contexts of its use and *selected* for representation in every one of those contexts. Automatic contents are those that show up across *every* use of some given concept. And this makes sense if the distinction between automatic and contextually cued content is exhaustive: automatic content is content represented in not only some of a concept's uses, but in all its uses, across all contexts, whether or not the particular context calls for its representation.

So, is being a *good companion* represented in all contexts of the use of DOG? Is *produce* represented across all contexts of the use of FRUIT? As far as I can tell, the answer in both cases, is *no*. As I've already suggested, some people consider dogs to be *dangerous* rather than good companions. These include people who might previously have considered dogs to be safe until a

traumatizing attack. And as we saw in the discussion of the botanical conceptions of fruit, botanists standardly include various inedibles, such as *winter berries*, under their botanical conception of fruit. This is incompatible with the idea that their use of the concept automatically represents fruit in such culinary terms as *produce*.

Indeed, in the case of identifying automatic contents for FRUIT, *Invariantism* seems to rule out just about every content I discussed in the previous Section: the concept can't include the knowledge that fruit are *seeded plant matter, developed from plant ovaries* (children don't represent that, at least while they're children), it can't include that they are *edible, sweet, juicy produce* (botanists don't represent that, at least while they're botanizing), it can't include reference to *apples, pears*, and *strawberries* (these aren't included in the botanical conception that refers only to the "true fruits"), it can't include reference to *legumes* or *nuts* (only some culinary representations include those), and it can't include reference to *pumpkins* or *olives* (people sometimes deny that they're fruits). As far as I can tell, the only content ignored, denied, or excluded by none of the participants and peoples I've discussed so far is the intentional content *grapes*, but I see no indication that grapes are selected for representation in *all* the contexts of the concept's use; indeed, I noted evidence that its representation can be encouraged, i.e. *cued by contexts*, that prime the idea that some fruits can be made into wine. I'm thus hard pressed to identify *any* automatic content for the concept FRUIT.

Now, it is open to *Invariantism* to make the claim that there are multiple distinct concepts of fruit at work here, and I'll return to this suggestion below, but first I want to note the problem that confronts *Invariantism* in the case that there really is just one shared concept FRUIT being employed across all these different contexts. If there is just one concept FRUIT and *Invariantism* is right, then there is little if any content actually represented by the concept FRUIT itself across

all the different contexts of its use and all or even most of the contents associated with its many varied usages are actually external to the concept itself. And if this is the case, then there appears to be little role for the concept itself to do in explaining our cognitive behavior with respect to fruit. Indeed, I see no reason to even assume that someone possesses the concept on the basis of their cognitive behaviors regarding fruit, no matter how sophisticated, if those behaviors show only evidence of contents that don't belong, as the proponent of *Invariantism* would have it, to the concept FRUIT. The problem for *Invariantism* is thus that it must give up the idea that concepts are to play an explanatory role in our cognition, at least in cases of highly variable contents, like those that are evident in the varied uses of the concept FRUIT.

Indeed, even in cases of less variable contents, the scope of possible cognitive import is diminished under the assumption of *Invariantism*. Suppose Machery is right that DOG automatically represents the knowledge that they are *four legged furry animals*. This is of limited cognitive import absent further so-called "background" knowledge about dogs. For consider, this content alone does not even allow one the ability to distinguish dogs from most other mammals. Things can be improved of course if we add in the knowledge that dogs *bark*, but we still wouldn't be able to distinguish dogs from other barking mammals: wolves, dingoes, foxes, seals, gorillas, several sorts of antelope, and more are indistinguishable from dogs if we employ only the automatic contents represented by Machery's analysis of DOG.

I think this result of explanatory poverty is an unhappy consequence for *Invariantism* and one that is not adequately appreciated in the literature. Previous critics have noted that the view strips concepts of too many plausible contents, but few appreciate just how empty and inert some of our most ordinary and everyday concepts are rendered as a result. Better then, one might think, to explore the idea that some such cases are to be understood as involving multiple *distinct* concepts, each representing their own automatic contents across all the contexts of their use. But this suggestion has its own rather unhappy implications too.

4.4.2.2 The Problem of Empirical Invalidity

As I've already mentioned, one of the main advantages of adopting *Invariantism* is that it nicely accommodates the idea that concepts can both persist over time, enduring through changes and variations in their associated auxiliary contents, and that concepts can be shared, despite large differences, divergences, and disagreements in auxiliary contents. This advantage is seriously diminished, however, if different people and the same people across different stages of their development, education, and training are employing distinct concepts. This is one shortcoming for the suggestion that what we'd think is a single concept, e.g. FRUIT or TREE or TIME, is in fact many different concepts of fruit, trees, and time. But there is a much more serious problem with the suggestion of splitting up our apparently persisting and shared concepts.

Recall the motivation for the move to split. Contents are automatic only if they are represented across all the contexts of a concept's use. But a great many, if not most, candidate contents are represented only on some occasions, in only some contexts, as is evidenced by cases such as that of FRUIT. The first (highly unpalatable) option is that our concepts just don't represent all that much content. Indeed, in some cases, our concepts are more or less *empty*. The second (apparently much happier) option is that different concepts are being employed across these different contexts. This allows for much richer contents: KID-FRUIT may include *apple* and *edible* as automatic contents, ADULT-NOVICE-FRUIT may include automatic representation of *plant matter* and *seeded*, CULINARY-FRUIT may automatically include *juicy produce*, BOTANICALLY-TRUE-FRUIT may include automatic representation of *grapes* and
pomegranates, while excluding *apples* and *pears*, NIGERIAN-FRUIT could automatically include *avocadoes*, while AFGHANI-FRUIT automatically includes *nuts*, and so on. The move to split thus saves *Invariantism* from the problem of cognitive import. But that does not make the move legitimate.

My primary concern with the move is that it renders *Invariantism* empirically invalid. Recall that contents are automatic only if they are represented across all contexts. And content belongs to a concept only if they are automatic. And so contents belong to a concept only if they are represented across all contexts of its use. This idea can be put to the test: take some concept, and some uncontroversial part of its content, and check whether it is represented across a large variety of contexts. If it is, then we have some evidence that favors *Invariantism*. And if it doesn't, we have some evidence that counters it. But if it is always open to *Invariantism* to reinterpret the latter outcome as one involving more than one concept, i.e. different concepts being employed in different contexts, then the possibility of counterevidence disappears; Invariantism becomes unfalsifiable. This is quite a large problem for a view that is framed as one that is empirically motivated, one that is responsive to the empirical evidence provided by our cognitive behaviors. But *Invariantism*_must make this move, if it is to retain the idea that concepts explain cognition. Invariantism thus either fails to respect the idea that concept enable cognition or they fail to respect the idea that theories of concepts are responsible for empirical evidence concerning cognitive behavior. Time to see what other interpretations might be on offer.

4.5 Against Contextualism

4.5.1 Contextualism

In defending his version of *Invariantism*, Machery relies on the distinction between *context-independent* and *context-dependent properties* posited by (Barsalou L., 1982). The former are properties represented by a concept across all (or most) of its various uses while the latter are properties represented only on occasion, according to contextual cues. This is the precursor to Machery's distinction between automatic content and contextually cued content. And yet Barsalou has subsequently moved to abandon the distinction, eventually suggesting that no property is wholly context-independently represented by a concept and instead that all properties are cued to some extent by context for representation by concepts.⁸⁷

This shift has motivated the adoption of a contextualist alternative to *Invariantism*, according to which all representational contents are contextually cued. Such *Contextualism* has come to represent the main alternative to the default adherence in the concepts literature to some form of concept invariance and stability.

Many of the most recent versions of such *Contextualism* have proposed a particularly radical picture according to which concepts themselves are contextually constructed. According to these accounts, concepts are merely temporary entities, existing only in working memory, representing only those contents that are contextually cued, exclusively so as to achieve

⁸⁷ (Casasanto & Lupyan, 2015, pp. 549, fn. 2), for example, report that, "In personal communication with D.C. [Daniel Casasanto] (October 20, 2012), L. Barsalou agreed that the CI-CD distinction [i.e. between *context-independent* and *context-dependent properties*] has outlived its usefulness." But Barsalou has also explicitly abandoned the idea that contents are represented absent contextual cuing in print. See, for example, discussion in (Barsalou L. , 2003), (Barsalou L. , 2005), (Yeh & Barsalou, 2006), and (Barsalou, Wilson, & Hasenkamp, 2010).

contextually given cognitive tasks. On this account, concepts aren't stored in long-term memory; only the existing bodies of knowledge from which they are constructed are stored for re-use across different circumstances of higher cognition. For example,

I use the term *concept* to refer *only* to temporarily constructed representations in working memory; *concept* will *never* refer to information in long-term memory. Instead a concept is simply a particular individual's *conception* of a category on a particular occasion. (Barsalou L., 1989), p. 93.

Concepts are situational-dependent mental entities. They are composed of semantic features which are flexibly recruited from distributed, yet localized, semantic maps in modality-specific brain regions depending on contextual constraints. (Hoenig, Sim, Bochev, Herrnberger, & Kiefer, 2008, p. 1799)

Rather than a process of accessing a preformed package of knowledge, instantiating a concept is always a process of activating an ad hoc network of stored information in response to cues in context. (Casasanto & Lupyan, 2015, p. 546)

On such an account then, there is no one concept FRUIT. There are only many concepts of fruit, each representing an idiosyncratic body of epistemic and intentional contents as cued by the particular context in which each is cued for use, none of which are necessarily used again, or stored in long-term memory for potential future re-use. That is, each time someone performs some cognitive task with respect to fruit, a different body of characteristics may be attributed to a different category of objects. No content is guaranteed to recur across different circumstances of cognizing fruit. More drastically still, no concept persists, and no concept is shared:

On this proposal, which we call the *ad hoc cognition* (AHC) framework, all words are infinitely polysemous, all communication is "good enough," and no idea is ever the same twice. (Casasanto & Lupyan, 2015, p. 543).

Theorists who endorse *Contextualism* frame their view as a direct opposition to *Invariantism*, insofar as they deny all invariance, insofar as they posit no distinction between invariant and auxiliary content. And yet, their explanation of content variation ultimately treads a strikingly similar path.

According to *Invariantism*, recall, content variation occurs as a result of either some difference in auxiliary content, coupled with no difference in invariant content, or some difference in invariant content, but belonging to distinct concepts. But notice now that *Contextualism* must explain content variation in much the same way, by always endorsing the latter disjunct. For it holds that every cognitive act employs a new concept, i.e. distinct from all those that have been used before, and each new concept may represent contents different from those represented by those used before. And so, content variation between different circumstances must be the result of difference between distinct concepts.

Indeed, the similarity goes further still. *Contextualism* holds that concepts each represent exactly one body of content. Concepts are thus actually trivially invariant across different contexts, since concepts are not re-used across contexts and thus cannot vary in content across them. And so, paradoxically, *Contextualism* is committed to the claim that each concept represents exactly one body of invariant content and it is the difference between those invariant contents, represented by distinct concepts, that explain content variation. As much as *Invariantism* then, *Contextualism* adheres to the basic tenet of *Concept Stability*.

Despite this overlap, *Contextualism* offers at least one distinct advantage over *Invariantism*, which is that concepts are placed front and center in explaining our cognitive behaviors, including the variations that can be observed in our cognitive behaviors. Every cognitive act is enabled by a concept that represents content tailored to that exact context. And every apparent content, revealed by how we think, talk, and act, can be attributed to that concept. *Contextualism* can even account for the apparent stabilities in content that appear to favor *Invariantism*, for similarities and overlaps of contents can be explained by an appeal to similarity and overlaps of context. Nevertheless, the view is not without its own unhappy shortcomings.

4.5.2 The Problem(s) with *Contextualism*

As might already be evident, a significant disadvantage for *Contextualism* is that is gives up the idea that concepts are the sorts of things that can persist over time and be shared between people. This includes the possibility of imparting a concept to a student, workshopping a concept with a collaborator, tracing a concept's cultural history, and so on. If we take *Contextualism* at its word, there literally is no such thing as remembering, communicating, learning, contesting, changing, explicating, or improving such things as concepts. Concepts are not the sorts of things that show up in our long term semantic memory, or occupy a place in our enduring cognitive architecture.

Now, this consequence might already be sufficient for some to give up on the idea of *Contextualism*. Indeed, Fodor (for one) actually states the possibility of cross-contextual, cross-disciplinary, cross-cultural, and long-term historical sharing of concepts as a nearly necessary condition on any satisfactory theory of concepts:

Barring very pressing considerations to the contrary, it should turn out that people who live in very different cultures and/or at very different times (me and Aristotle, for example) both have the concept FOOD; and that some people who are possessed of very different amounts of mathematical sophistication (me and Einstein, for example) both have the concept TRIANGLE; and that people who have had very different kinds of learning experiences (me and Helen Keller, for example) both have the concept TREE; and that people with very different amounts of knowledge (me and a four-year-old, for example) both have the concept HOUSE. And so forth. Accordingly, if a theory or an experimental procedure distinguishes between my concept DOG and Aristotle's, or between my concept TRIANGLE and Einstein's, or between my concept TREE and Helen Keller's etc. that is a very strong prima facie reason to doubt that the theory has got it right about concept individuation or that the experimental procedure is really a measure of concept possession. (Fodor J., 1998, p. 29)

And I tend to agree with Fodor here, but it must be conceded that *Contextualism* has some non-trivial reasons to claim that it has actually offered us some "very pressing considerations to the contrary." As I see it, only *Contextualism* has (so far) provided a sufficiently serious accounting of the extent to which contextual variation actually shows up across concept use and cognition more broadly. The difference that context can make is radical and often drastically underestimated by even those theorists who take the experimental, sociological, and historical evidence of concept use seriously. Moreover, the idea that cognitive science can depart, sometimes very radically, from the folk psychology from which it takes its departure, is well-motivated in many corners of empirical psychology. Our folk notions of mental life can and often have been overruled by empirical findings. The same may well be true of the notion of a concept as something that can persist and be shared. *Contextualism* may well be on the right track here. Nevertheless, I don't think the evidence forces us to go all the way down the *Contextualist* road.

The "considerations to the contrary" are pressing only if there are no good alternative interpretations of the evidence supporting *Contextualism*. I've already shown why I think that

Invariantism is too problematic to be a good contender. But I've also already suggested a third option, i.e. *Concept Variability*, according to which concepts store a variable set of contents, each of which may be updated, adjusted, or replaced as an individual gains greater conceptual sophistication, and each of which may, but needn't, be selected for representation in any given occasion of the concept's use in producing some cognitive act. On this view, there are no necessarily invariant contents stored, no contents privileged for automatic representation, each content may be selectively cued by and for particular contexts, and yet concepts can *persist* through such changes, and nothing prevents them from being *shared*. We can make sense of the data that motivates *Contextualism* without going the unnecessarily radical *Contextualist* route. We can preserve the idea that concepts can persist, that they can be shared, that their contents sometimes appear stable, while also taking care of the evidence that their contents can vary across contexts. But we can only do so by rejecting *Concept Stability* and adopting *Concept Variability* in its stead.

4.6 Against Concept Stability

The primary advantage of *Invariantism* is that it accounts well for the idea that concepts can persist and be shared across different contexts of use, degrees and directions of conceptual sophistication and disciplinary expertise, as well cultural divergences. Its primary disadvantage is that it struggles to provide a satisfactory explanation of content variations across these different contexts and cognizers. In particular, in positing invariant concepts, *Invariantism* must either give up the idea of that concepts enable cognition, or risk empirical invalidity. The situation with *Contextualism* is exactly the reverse: it scores well on explaining content variation, but fails to retain the idea that concepts can persist across contexts and be shared across cognizers. It only

captures the true extent of content variation by abandoning concept persistence. The problem in each case is an adherence to *Concept Stability*, the idea that concepts represent a stable set of contents. If we reject this traditional idea, and adopt *Concept Variability* in its stead however, we are free to enjoy both advantages, while avoiding both the disadvantages; we can capture both content variation and concept persistence.

According to *Concept Variability*, concepts can persist across changes in their contents, they can be shared by people with different perspectives, experiences, expertise, and training, and they can represent different selections of their stored contents according to different contextual cues, thereby representing their contents in order to enable a large variety of different cognitive behaviors.

The abductive argument for *Concept Variability* that I have offered here may well be sufficient to convince those already dissatisfied with the accounts provided by *Invariantism* and *Contextualism*. But I think an even stronger case can be made in favor of this alternative to *Concept Stability*, one which relies even more directly on the idea that concepts are cognition-enablers, those cognitive units of central importance in explaining our ability to engage in higher cognitive behavior. It is this stronger argument that I turn to next.

5.0 The Cognitive Necessity of Concept Variability

5.1 Introduction

In the previous Chapter I made an abductive argument in favor of *Concept Variability*, by showing that the denial of *Concept Stability* allows us to make better sense of both concept persistence and content variation in concept use. In this Chapter I make a deductive argument that concepts admit variation in their representational contents, showing that content variability is a necessary condition for concepts to enable cognition. In many respects, the basic components of this argument have already been put in place by my remarks and claims in the previous Chapters. They all come together here, to complete my case for *Concept Variability*.

The core argument proceeds as follows: *Premise 1*: Concepts enable cognition, by representing contents. *Premise 2*: Concepts can enable cognition only if they admit of variation in their representational contents. *Conclusion*: Concepts admit of variation in their representational contents.

As I've already indicated in the previous Chapter, *Premise 1* is not in much dispute. The idea that concepts lie at the center of our ability to think and reason, and that they do so by supplying the contents of thought, is probably the only claim generally accepted across the theorizing of concepts in philosophy, psychology, and cognitive science in general. It is by endorsing *Premise 2* that I break from the orthodoxy. According to most standard theories, concepts are to be individuated precisely by their invariant contents, which thus cannot vary without a concept losing its very identity. In contrast to such accounts, I argue that cognition *requires* concepts to admit of variation in their contents; it is only *because* concepts admit of

variation in their representational contents that they can enable cognition. This of course, is in stark contrast to the view defended by proponents of *Concept Stability*, according to which concepts must be *stable* if they are to enable cognition. As Barsalou and Camp put it,

These intuitions of stability are often compelling, and it is sometimes hard to imagine how we could communicate or perform other intelligent behaviors without *stable* knowledge structures [i.e. concepts]. (Barsalou L., 1989, p. 76; my emphasis)

[...] concepts are not essentially tied to any particular attitude: conceptual thought enables thinkers to enter into multiple mental states [...] with the represented content remaining *constant* across those attitudinal changes. [...] But for concepts to be context-independent in this sense, they must also be cross-contextually *stable*: the same concept must be able to be redeployed on different occasions and in different applications with a *common* representational import. (Camp, 2015, p. 593; my emphasis)

I argue here for the exact opposite: cognition requires variability, not stability, in content. I proceed as follows. I begin in §5.2 by recapitulating what I mean by such terms as "concept" and their "content" in order to articulate more precisely the content of *Premise 1* and the thesis of *Concept Variability*, before illustrating its intuitive appeal with my central example of FISH. In §5.3 I turn to my defense of *Premise 2*, relying again on the empirical evidence that indicates the role of content variation in cognition, this time showing that the idea that concepts are cognition enablers *requires* that they admit of variation in their contents. I respond to several tempting objections along the way before closing in §5.4.

5.2 Concepts and their Many Contents

My aim in this Section is to re-articulate the thesis of *Concept Variability*. I do so by first stating again some of my central assumptions concerning concepts and their contents, which together constitute my interpretation of *Premise 1*, before exemplifying the thesis with the final example of FISH. I also use this Section to respond to a particularly tempting objection to my thesis, according to which my chosen example should be understood as representing exactly one particular body of contents rather than, as I hold, a variety of contents across various cognizers and their various acts of cognition.

5.2.1 Concept Variability Articulated

My central assumption throughout the last few Chapters has been that *concepts enable cognition*. Recall that by this I mean what many theorists of concepts mean: concepts are *explanatory posits*, posited to explain the ability of cognitive agents to perform several sorts of higher cognitive acts. Most paradigmatically, concepts are invoked to explain *categorization*, the ability to sort objects into different categories, and *inference*, the ability to project information across the members of a category. It is by appeal to the concept FISH, for example, that we can explain a cognizer's ability to sort between fish and non-fish. And it is by appeal to FISH that we can explain a cognizer's ability to project the properties of some particular fish to fish more generally, and thus their ability to infer from something's being a fish that it has certain fish properties.

More specifically, I join many theorists in assuming that concepts enable cognition by *representing contents*, i.e. it is by representing their contents that concepts form part of the

explanatory chain that terminates in the performance of some particular cognitive act. It is because FISH represents contents that concern fish that it enables cognizers to think and reason about fish. But there are at least two things that I've come to mean by "representation" as well as at least two things referred to by the use of "content".

By "representation" we might mean *storage* or we might mean *activation*. That is, we might think of concepts as enabling cognition by *storing* their contents, or we might mean that they enable cognition by *activating* their contents. As became clear in the last Chapter, I mean both. Indeed, I think the latter presupposes the former; concepts enable cognition by activating their stored contents.⁸⁸ The distinction is worth drawing however, insofar as concepts can, on my view, activate a mere *selection* of its stored contents in the course of enabling any particular act of cognition.

By a concept's "content" we might mean its *intentional content* or we might mean its *epistemic content*. Again, I mean to refer to both. Recall that intentional content, on my view, is constituted by a concept's *referents*, i.e. those categories of entities (broadly construed) to which a concept refers. Depending on the particular concept in question, such entities might be objects, properties, or events, real or unreal, singular or general, abstract or concrete, empirical or not. In the case of FISH, its intentional content is *fish*. Epistemic content, on my view, is constituted by those *bodies of general knowledge* a concept attributes to its objects. And in this case, what is

⁸⁸ I do not mean by this to deny that we also sometimes *derive* "ad hoc" contents in the course of using a concept. These are contents that are active without having been activated from storage, in virtue of having been derived *in situ* for some particular purpose (see, for example, discussion in (Barsalou L., 1983), (Carston, 2002), and (Allott & Textor, 2012). Such contents also have an important role to play in cognition, but I will simplify matters here by assuming that their role is subsequent to the work of stored and activated content, and thus downstream of my interests here.

included among the epistemic contents of any given concept, including the concept FISH, is, as we've seen, a matter of much dispute.

According to *Definitionalism*, the epistemic content of a concept is exhausted by its *definition*, i.e. an articulation of those properties that are individually necessary and jointly sufficient for any object to be among the referents of that concept. But, more contemporary theories, including that of *Essentialism* and other forms of *Hybridism*, as well as different sorts of *Typicalism*, have mostly given up on such a restrictive account of a concept's epistemic contents, holding instead that concepts can represent characteristics that are merely *typical* of their objects, And as we've seen, these characteristics might be articulated by, among other things, a *statistical prototype* of a concept's objects, or by an *essentialist theory* that connects the characteristics of those objects, or by some combination of these structures and others.

I spent the first two Chapters arguing that even these more contemporary accounts are in some respect too restrictive, and that a concept's epistemic contents should include a wider variety of contents that characterize its objects. Accordingly I have come to think of a concept's epistemic contents as articulated by the *generic characteristics* of a concept's objects, i.e. those characteristics that are generally present among a concept's objects, but which may admit of a number of exceptions in any particular case, while leaving open the conditions under which a characteristic is generic for any given category of objects.⁸⁹ Such an account allows the cognitive contents of FISH to include such typical fish characteristics as *living underwater, breathing oxygen through gills, being cold-blooded, having finned limbs*, etc. even though, as we'll see in more detail

⁸⁹ That is, the genericity of a characteristic could be grounded in its being definitional, essential, prototypical, stereotypical, explanatory, or something else entirely. My own view, on which I have endeavored to stay neutral, but which might be evident from my choice in examples, is that characteristic genericity is *domain-relative*. See my (Visser, 2023), in which I develop and defend this view, according to which a characteristic is generic for some category of objects in virtue of being invoked as an *explanans* in any of those domains that also invoke the category.

below, not all fish exhibit all these characteristics, and not all fish-involving disciplines and domains take particular interest in all such fishy properties.

Having restated my two guiding assumptions about concepts and their contents, I can now state more concretely the thesis of *Concept Variability*. The thesis holds that concepts admit of variation in their representational contents. By this I mean that concepts admit of variation in both the intentional contents and the epistemic contents that they store and activate in order to produce particular acts of cognition, such as categorization and inference. Concepts can be used to represent different categories of objects and different characterizations of those objects across different acts of cognition, activating different selections of these contents on any given occasion of its use and storing different combinations of these contents across different cognizers, as well as the same cognizers over time. I spent the previous Chapter offering an abductive argument for this claim. The deductive argument I offer here turns on the claim that it is only *because* concepts admit of such variations in their representational concepts that they can be invoked to explain those acts of cognition. Cognition is made possible by variable concepts.

I've sometimes also articulated the thesis of *Concept Variability* is by appeal to the notion of a "conception". Aiming to capture a sense in which different cognizers can conceive of the same subject matter in different ways, some theorists have drawn a distinction between a concept, like FISH, and its many conceptions, as might be exemplified by an ichthyologist's narrower biological understanding of fish in contrast to a chef's broader culinary understanding.⁹⁰ This is a useful way of articulating the idea of content variation as can be exhibited by a single concept, as long as we do not presuppose, with many of these theorists, that cognizers possess no more than one conception of any given concept: an ichthyologist might employ a biological conception of FISH

⁹⁰ See, for example, (Rey, 1985), (Millikan, 2005), (Rey, 2010), and (Lalumera, 2014).

at work, but a culinary conception at dinner, while a chef might employ a culinary conception at work, but a biological conception while watching a nature documentary. Nothing necessities that cognizers don't store a variety of conceptions for a single concept, any combination of which might be used to activate a selection of contents in order to enable any particular act of cognition.

5.2.2 Concept Variability Exemplified

As has become evident, my favorite examples of variable concepts are those that admit of variation as a result of being simultaneously active in several different interrelated disciplines or domains, while also being identifiable across many of the world's cultures. My central example of such a concept here is the widely familiar concept FISH, which resembles the last Chapter's FRUIT in several ways.⁹¹ Most of us know that fish are underwater organisms, typically finned, gilled, and covered in scales. Many of us also know that fish are the sorts of things that ichthyologists investigate, fishers fish, and pescatarians prefer to eat, if they're eating meat. Fish are also among the things that climate activists and marine environmentalists aim to protect and the extraction of and interaction with which legislators across the world aim to control by implementing and enforcing their various fishing regulations and laws. And yet, across these various domains, conceptions of fish can vary.

⁹¹ Although my focus here is on the interconnected conceptions that can be expressed in English with the word "fish", many of my claims hold for a variety of other contemporary languages too. Examples of which I am aware include not only the closely related Afrikaans "vis", Dutch "vis", and German "Fisch", but also Slovenian "riba", Portuguese "peixe", Spanish "pescado", Mandarin "鱼", Arabic "سياك", Urdu "مجهلى", and Persian "مجهلى". It is true also of the American Sign Language and British Sign Language signs for fish. In each of these cases I've relied on the testimony of fluent speakers, but their verdicts can be confirmed by an investigation into what sorts of objects and characterizations are included (and presupposed) in fish textbooks, fish recipes, fish laws, and other cultural documents that concern fish in each linguistic tradition.



Figure 3: Various overlapping referents of FISH

Consider first some variations in the intentional contents of FISH (See Fig. 3 above).⁹² What counts as a fish in biological domains tends to be narrower than what counts as a fish in different culinary cultures, which itself tends to be narrower than what usually counts as a fish across various legal jurisdictions.⁹³ We also know that what counts as a fish today is not what counted as one in these domains in the past. Assuming that these different referents all belong to

⁹² This figure is necessarily partial and schematic. Many more candidate referents can be added. And many of the boundaries represented here might be rather vague as well as inter-culturally and inter-personally variable. Many more conceptions of fish have existed and continue to exist in many different domains today.

⁹³ Culinary conceptions typically include *shellfish*, among others. Taking a wide international view, the U.N.'s FOA holds that, in the culinary sense, what is called "fish" includes "finfishes, crustaceans, molluscs, amphibians, freshwater turtles and other aquatic animals (such as sea cucumbers, sea urchins, sea squirts and edible jellyfish) produced for the intended use as food for human consumption." (FOA, 2014). Legal conceptions are typically even wider. According to U.S. law, for example, "*Fish* means: (1) When used as a noun, means any finfish, mollusk, crustacean, or parts thereof, and all other forms of marine animal and plant life other than marine mammals and birds." (50 CFR §600.10). Meanwhile, the legal definition in India excludes plants while possibly including *aquatic birds*: "Fish' means an animal whether living or dead of a species that throughout its life cycle usually lives—(i) in water (whether fresh water or salt water); or (ii) in or on foreshores; or (iii) in or on land under water, which includes the spat, spawn and eggs of fish." (Act No. XVI of 2018, §2).

the single concept FISH, its category boundaries are not only variable between different interrelated domains, but have also changed over time. As I see it then, the objects that FISH represents has not only changed over time, but also continues to vary across different contemporary fish-involving domains, with individual cognizers capable of thinking of different fish across different uses of the concept FISH.

Consider next the epistemic variations that can accompany these intentional variations. In biology, the predominant interest is in describing the biological properties of fish, those that characterize their morphology, ontogeny, phylogeny, and ecology, but little emphasis is placed on their culinary potentials or their cultural and legal significances. In culinary domains, there is much greater interest in the culinary properties of fish—what they taste like, what sorts of nutrients, allergens, or toxins they contain, which sorts of fish pair well with which sorts of other culinary flavors, their religious relevancies, their ritualist roles, etc. And in legal domains, the emphasis is more on how fish are caught, displaced, or threatened by human activity in and around the water as well as in their distribution through our many trade networks for the purposes of human exploitation and consumption.

Such characterizations have also changed over time.⁹⁴ We know much more about fish today, especially about their inner lives, than we did even a few decades ago.⁹⁵ And as our technologies of fish investigation, extraction, transportation, consumption, and preservation have changed, as well as our ideas about their place in the global economy and in our moral universe, so too have our various conceptions of them. To the extent that these epistemic variations all

⁹⁴ See especially (Simoons, 1994) who provides a highly detailed account of the different ways fish have been conceptualized, both positively and negatively, across world history in different societies and the different cultural roles they have played in those societies over time.

⁹⁵ See, for example, (Balcombe, 2017), who argues forcefully for the overwhelming evidence available today for the cognitive and social complexity of several different sorts of fish, or as he prefers to refer to them, "fishes".

belong to the same concept FISH, it is thus a concept that admits of a great many different characterizations of fish.

5.2.2.1 Interim Objection 1: Don't Fish Constitute a Natural Kind?

Before I go any further, it is worth noting that some readers may be tempted to dismiss my use of the example of FISH on the grounds that fish constitute a *natural kind*, i.e. a category whose boundaries are settled in some important metaphysical sense by nature itself, independently of any human interests or ideas. Accordingly, such readers might be tempted to claim that there is actually some uniquely privileged conception of fish, constituting the "true" content of FISH, i.e. whichever conception has as its intentional content the natural kind of fish (or perhaps that conception that has as its epistemic content the wholly accurate characterization of the objects that belong to the natural kind of fish). Such readers may thus hope to reject *Concept Variability*, holding that while there may be many different conceptions of fish across many different domains, applying to different categories of objects and characterizing those objects in different ways, all but one are, ultimately, mere *misconceptions* of fish.

My first response to this objection is that relatively few (if any) of our concepts are uniquely applicable to a single natural kind to the exclusion of all other nearby overlapping categories, if only because many of our concepts are not of natural kinds. We regularly conceive of categories of natural objects that belong to different natural kinds—consider *weeds*, or *bugs*.⁹⁶ And we regularly conceive of categories that aren't, at least in any straightforward sense, determined by the natural world, independently of any human interest or concerns—think *weddings*, or *blogs*. So

⁹⁶ See (Weiskopf D., 2020) for a nice discussion and several illuminating examples of such "anthropic" concepts.

it cannot be that for each of our concepts there exists exactly one natural category that it has as its "true" intentional content (or some perfect characterization of that category that it has as its "true" epistemic content). The response thus cannot generalize, even if it does apply to the case of FISH.

My second response is that the objection does not apply even to the case of FISH, for fish do not constitute any particular natural kind. Indeed, the domain that most plausibly has the right to claim its conceptions of fish as somehow privileged by the natural world explicitly holds back from naturalizing its subject matter in this way. As one prominent textbook of ichthyology notes:

To most biologists, the term 'fish' is not so much a taxonomic ranking as a convenient description for aquatic organisms as diverse as hagfishes, lampreys, sharks, rays, lungfishes, sturgeons, gars, and advanced ray-finned fishes. (Helfman, Collette, Facey, & Bowen, 2009, p. 3)

Another textbook warns:

"Fishes" is not a monophyletic group (i.e., a group made up of an ancestor and all of its descendants) because the tetrapods, which share a common vertebrate ancestor with fishes, are excluded. Thus "fish" typically refers to any vertebrate that is not a tetrapod. Fishes (usually) live in water, (usually) obtain oxygen through gills, are (usually) ectothermic (i.e., cold blooded), and (usually) have limbs in the form of fins. Naturally, there are exceptions to each of these rules. Some fishes spend time out of the water, some breathe air, some are endothermic (i.e., warm blooded), and some have no limbs at all. (Hastings, Walker, & Galland, 2014, pp. xvii–xviii)

That is, even among ichthyologists there is no general presumption that there exists some naturally privileged category of organism that constitutes the subject matter of ichthyology. There is no unique shared morphology nor any uniquely shared natural history.⁹⁷ And so there is no particular natural kind known to or posited by natural science to which the concept of fish is to apply to the exclusion of all other similar categories. Nothing in the natural world can thus privilege any one biological conception of fish over any other, or indeed over any culinary or legal conception of fish, across all contexts.⁹⁸ They are all simply different conceptions of fish, active in different interrelated domains that include an enduring interest in the many interrelated theoretical and practical potentials of fish.

There are many such things as fish, and many such things as a conception of fish. On my view, these conceptions all belong to one concept, the concept FISH, a concept that admits of various interrelated conceptions, in different interrelated domains, representing different interrelated contents, in order to enable different interrelated cognitions. But in order to fully defend this claim, we need to take a closer look at the role of concepts in enabling cognition, which, on my view, requires that one and the same concept admit of different contents across different acts of cognition, persisting through variations and changes in a single concept's many representational contents.

⁹⁷ There is thus also certainly no fish-essence of fish-definition to be discovered by natural science.

⁹⁸ Indeed, I follow (Dupré, 1993) and (Dupré, 1999) in holding that most categorization decisions both within and beyond the sciences are *functional* groupings, well-motivated only from the perspective of domain-specific interests. This means that insofar as different interests might motivate different groupings, many different overlapping categories can qualify as one among many different more or less natural kinds, both within the same domains as well as across different ones. And so, if there are any natural categories of fish, or indeed of any other natural entities, they are *plural*, with no one being more privileged than any other across all contexts. See also again (Taylor & Vickers, 2017) for an expansive list of examples of scientific concepts that admit of several different theoretical articulations, or as I would put it, admit of several different intentional and epistemic contents, thus constituting several different scientifically legitimate conceptions of the same concept. See also again the allied work of (Wilson, 2006)) on the "wandering" concepts of physics and their use in various applied sciences, and the especially illuminating work of (Lakatos, 1976) on the "growth" of mathematical concepts.

5.3 The Enabling Role of Content Variation in Cognition

Having articulated *Concept Variability* and some of the assumptions that frame it, my aim in this Section is to defend *Premise 2*, according to which concepts can enable cognition only if they admit of variation in their representational contents. My focus here is thus on showcasing some of the ways that cognition *depends* on concepts allowing the storage and activation of different intentional and epistemic contents across different acts of cognition. Since I have already motivated *Premise 1*, according to which concepts enable cognition by representing their contents, my arguments in this Section will complete my defense of *Concept Variability*.

The argument proceeds in four stages. I start by defending the narrower claim that categorization depends on concepts admitting variation in the activation of their epistemic contents. I then generalize the claim in three steps. First, I show that the dependence on variable activation of epistemic content generalizes beyond categorization to cognition in general. Second, I show that the dependence on variable activation generalizes beyond epistemic content to intentional content and thus to a concept's representational content in general. And third, I show that the dependence on variation generalizes beyond activation to storage. The result of these generalizations is *Premise 2*, which holds that cognition depends on concepts admitting variation in both the activation and storage of both their epistemic and intentional contents.

I bolster these arguments by responding to some particularly tempting objections along the way, closing the Section by responding once again to the objection according to which cognition might be enabled by a variation *between* concepts, rather than any variation *within* them, and thus that the things one might call "conceptions" can be recast as "concepts" in their own right without loss.

5.3.1 Categorization Requires Variable Activation of Epistemic Content

We can begin with a simple case of *categorization*, the cognitive ability to sort objects into different categories. The role of concepts here is to represent the contents that serve as the basis for *recognizing* a concept's objects, allowing them to be categorized accordingly. It is by representing contents that concern fish that FISH allows cognizers to recognize fish, and thus allows us to sort between things that are fish and things that are not. And yet, in order to enable different acts of categorization, the concept must admit of variable activation, allowing different contents to be represented across different acts of categorization.

One important reason for this is that much of the epistemic contents represented by a concept fail to extend across all of its intentional contents. As I argued at length in Chapter 2, concepts regularly represent merely typical, i.e. exception-admitting, characteristics of their objects. But this means that the recognition of a concept's different objects will sometimes have to depend on the representation of different characteristics, to the extent that different objects lack different characteristics represented by the concept to which they all belong. With respect to FISH, for example, one cannot rely on the characterization of fish that they *live and breathe underwater* in order to recognize an *African lungfish* as a fish, since they can live in dried mud for several years, using their lungs to obtain oxygen from the air. Similarly, one cannot rely on the characterization that fish have *finned limbs* in order to recognize a *hagfish* as a fish, since they have none. Indeed, the point extends to all contingent properties: fish aren't always *in the water*, they aren't always *alive*, they don't always have their *heads*, *fins*, *gills*, and *scales* still attached to their bodies, and so on. The ability to recognize the different objects represented by a concept will thus depend on the ability to activate different characterizations of those objects across different

circumstances. No single characteristic will be of use across all the categorizations a concept serves to produce if that characteristic is only sometimes present.

A second important reason that different acts of recognition must sometimes depend on different epistemic contents across different contexts is that cognizers are not always in the position to recognize the characteristics of objects, even when they are present. We cannot, for example, rely on perceptual properties—colors, tastes, smells, etc.—when sensory conditions are poor or our sensory capacities are impaired: we cannot rely on the typical texture of a fish if we cannot touch it, and we cannot rely on its typical visual appearance if we cannot see it. But, by the same token, we might sometimes have to make do with only the superficial properties of some object if our ability to investigate or discover its underlying characteristics are limited: we cannot rely on the characteristic thermal properties of fish if we are not in a position to assess or understand the mechanisms that control a particular fish's body-temperature, and we cannot rely on whatever genetic knowledge we might possess about fish, if we are not in a position to assess a particular fish at a cellular level. That is, even if it is not always contingent whether or not the objects represented by some concept manifests some particular characteristic, it remains contingent whether a cognizer is in a position to discern it on any given occasion.

And so, for a wide variety of properties characteristic of a concept's objects, there are contexts where we are not in the position to discern that property while still being in the position to discern some others, either because the characteristic is absent in that particular case or because we are not in a position to notice or confirm its presence. Recognition thus depends on our ability to activate several different epistemic contents, each of which will only sometimes be of any use in recognizing some of a concept's objects as such an object. Different contexts thus call for the activation of different selections of a concept's epistemic contents even if we hold fixed the object to be recognized, the cognizer performing the recognition, and the concept that enables the recognition. Categorization requires concepts to admit of variable activation of their epistemic contents, because different selections of a concept's contents will explain how different acts of categorization were made possible by that concept.

5.3.2 Cognition Requires Variable Activation of Epistemic Content

The need for variability in content activation extends to other types of cognition. Indeed, the point holds *across* different types of cognition, for different types of cognition can depend on the representation of different sorts of contents. In this respect, it is constructive to compare the sorts of epistemic contents that tend to enable *categorization* and the sorts that tend to enable *inference*.⁹⁹

Categorization tends to depend on more *distinctive* characteristics, properties that are more common (and especially more commonly discernible) among a category of objects than among its contrast categories. To take a rather extreme example, it is not very helpful to rely on the characterization of fish that they are *existents* even though this is true, since this does little to distinguish a fish from a great many other things. The same can be said of such mundanities as the facts that fish *occupy time and space*, that they are *subject to gravity*, that they are *organisms*, perhaps even that they *live underwater*, for these are true for many of the categories of objects that stand in contrast to fish.¹⁰⁰

⁹⁹ See especially the discussion in (Markman & Ross, 2003), where the focus is on the contrast between the contents that enable categorization and the contents that enable predictive inference.

¹⁰⁰ As I emphasized in Chapter 3, it is not obvious what makes for a contrast class of any given category. Indeed, I think that whether or not a category counts as a contrast class for another can vary by context. In culinary contexts, for example, the contrast categories for fish might include, *meat*, *dairy*, *vegetables*, etc. But in biological contexts, the contrasts might be *mammals*, *birds*, *reptiles*, etc. If this is right, then which contents will count as distinctive of fish

The reverse is true for *inference*, the cognitive ability to project characteristics across a concept's referents. In this case, cognition tends to depend on characteristics that are *common* among a category of objects *regardless* of their prevalence among other things. Indeed, the more common, the better, for these contents produce the safest projections and predictions. From the fact that something is a fish, we can more safely infer the above mundanities than facts concerning less prevalent distinctivities of fish, e.g. that a given fish can *change its sex*, something which is very distinctive of fish, and yet something of which only a few hundred species of fish are capable, and then only under particular social and reproductive conditions. The presence of such a property can thus be useful in categorizing an object as a fish, but its presence cannot be safely inferred from the mere fact that something is a fish. Which epistemic contents are most usefully activated can thus vary across the different types of cognition a concept's contents serve to enable. Different contents will explain how different acts of cognition have been produced, with no one characteristic represented by a concept serving to explain every cognitive act enabled by that concept.

I think the same general point holds across all types of cognition enabled by concepts. That is, in each type of cognition in which concepts play an enabling role, I think we should expect variability not only between different acts of the same types of cognition, but variation between different types too. Each of a concept's epistemic contents will only sometimes be of any use in producing some act of cognition; no content will be usefully activated across every act of cognition. And so, if we are to invoke concepts in order to explain cognition, we must invoke different selections of their contents to explain different acts of cognition. In order to enable

will vary across these contexts, and so the categorization of fish will depend on different contents in different contexts. If this is right, then the contextual variability of contrast provides further support that different acts of cognition rely on the possibility of representing different contents.

cognition, concepts must admit of variation in which of their epistemic contents are activated across different acts of cognition.

5.3.3 Cognition Requires Variable Activation of *Representational* Content

My argument for variable activation is not complete until I've shown that the claim extends to a concept's *intentional* contents. So far, I've focused primarily on the variable activation of a concept's *epistemic* contents, those contents that characterize a concept's referents. But my view includes the claim that cognition requires variable activation in *intentional* contents too.

Consider again the case of *categorization*. Empirical studies have shown that cognizers do not always rely exclusively on a concept's epistemic content to categorize objects. They will sometimes depend on the representational activation of a concept's *exemplars*.¹⁰¹ Now, theorists aren't always clear on whether by "exemplar" they mean some particular *object* among the referents of a concept or rather some specific *subtype* of its referents; some would hold that it is *Nemo* that can serve as an exemplar of FISH, while others would hold instead (or perhaps in addition) that it is *clownfish*. Nevertheless, the evidence shows that cognizers sometimes rely on a representation of a selection of a concept's referents in order to perform categorization tasks rather than (or in addition to) a representation of their characteristic properties, and thus on a concept's *intentional* contents rather than its *epistemic* contents. But, *which* of a concept's referents will be of any use in any given act of categorization can be expected to vary in much the same way as its characterizations, for the activation of an exemplar can also be cognitively useful only relative to

¹⁰¹ As I mentioned in Chapter 3, this is the sort of evidence that has motivated the development of *Exemplarism*. See again overviews of this literature in (Murphy G., 2002) and (Machery E., 2009).

the contingencies of the objects to be recognized and the conditions under which a cognizer attempts the recognition. Activating a representation of some *clownfish* will only sometimes be of any use in attempting to recognize whether something is or isn't a referent of FISH; I suspect a representation of some clownfish will be of more use in categorizing, for example, other colorful little fish than *large grey sharks*, or *long dark eels*. If this is right, our concepts must admit of variable activation of their referents just as they must admit of variable activation of their characterizations, in order to enable different acts of categorization.

What about *inference*? In my discussion above, I focused on inferences that attribute a characteristic to some of a concept's referents on the basis of being included among its referents, e.g. inferring that an object can change its sex on the basis of knowing that it is a fish. But inferences can also involve projecting properties manifested by particular referents to a concept's referents more generally, e.g. inferring that fish in general are capable of *feeling pain* on the basis of evidence that some fish can hurt. The interesting thing to note here then, is that we can expect different sorts of characteristics to be projected across different categories of objects, each of which are among a concept's stored referents. That is, which fish we generalize over when projecting a property across the referents of FISH can vary across different sorts of categories of objects.

Consider again the different domains that broker in different conceptions of fish (see again Fig. 3). Part of what distinguishes these domains is that they admit different objects into their categories of fish and that they characterize those fish in different ways. What I'd like to draw attention to now is how these different categories can serve as the referents over which different inferences might project newly acquired information about the characteristics of fish.

Suppose I direct you to a recent finding concerning *stingrays* and *cichlids*; scientists recently found that these fish are capable of *basic arithmetic*. Popular media reported this result as "Fish can learn basic arithmetic" (Savitsky, 2022). The results have been projected across fish. But *which* fish? Most plausibly, insofar as the characteristic discovered is one that concerns the cognitive capacities of the studied fish, which is a biological property, the category to project the property over is a biological category of fish. But, what if the property concerned the legal status of fish? If you are told that you require a license to catch fish in some given jurisdiction, you better project that information across legal fish, whether or not they are included among biological or culinary fish. That is, different acts of inference may require a cognizer to project newly acquired information about a concept's objects over different categories of its referents.

The extent to which a cognizer projects information concerning fish over narrower categories in some cases but wider categories in others indicates their ability to vary between the referential categories activated by means of FISH. I conclude then that cognition requires variable activation of intentional contents across different acts of cognition in addition to requiring variable activation of a concept's epistemic contents.¹⁰²

5.3.4 Interim Objection 2: Why not *Plural* Activation?

I've argued that different acts of cognition are to be explained by the activation of different representational contents across different occasions of a concept's use, with no one content capable of serving to explain every act produced by the use of one and the same concept. Supposing that

¹⁰² Of course, the ability to shift between categories in this way remains consistent with my opponent's claim that we're activating different concepts, rather than different referents of one and the same concept. I'll address this alternative in §5.3.6.

this is right, this does not yet get us to the claim that cognition depends on content *variability*, for it remains possible that cognition proceeds via the activation of a concept's *every* content regardless of its contextual relevance. That is, it remains possible that whenever we perform any act of cognition concerning fish, we represent every fish and every characteristic of fish, every content we attach to the concept FISH, regardless of the utility of doing so in order to perform any particular act of cognition.

I think there are at least two reasons to doubt the plausibility of this possibility and to hold instead that cognition proceeds by *selectively alternating* between the representation of different contents for different cognitive acts, and thus in favor of *variability* in activation rather than mere *plurality* of activation.

The first reason is simple: representing its every content every time we attempt to use a concept in any act of cognition is highly *inefficient*. Indeed, it might be highly *confusing*, insofar as we cannot focus on all contents at once and some contents are actively counterproductive for the achievement of particular sorts of tasks. The speed and accuracy of categorization decisions for some objects may be undermined by the representation of characterizations that are not distinctive of the category to which it must be determined whether they belong. We might waste our time drawing inferences, expectations, or predictions that concern only a small portion of a concept's objects. And urgent projects might be delayed or derailed by the representation of alternate categories all at once or characterizations that recommend at least practically incompatible strategies.¹⁰³

¹⁰³ With respect to the latter, consider a recent legal decision: a court in California ruled—see, e.g. the report of (Sottile, 2022)—that *bees* are to be included under the legal category of fish for the purposes of the California Endangered Species Act, which applies to several sorts of animals, including fish, but which makes no mention of insects. As is well appreciated now, bees are extremely important to the health of many ecosystems, including the ecosystems of human agriculture, but they are also seriously endangered. Their protection could have been secured

The second reason that I think we should hold that cognition requires representational *variability* rather than mere representational *plurality* is that it makes better sense of the empirical literature which shows again and again that contextual factors influence which contents, both cognitive and referential, are represented in and across different acts of cognition. Recall the findings reported in Chapter 4: the literature includes evidence that context primes the activation of some characterizations instead of others as well as the activation of some exemplars instead of others, in addition to evidence that shows that cognizers are willing to apply their concepts to sometimes wider and sometimes narrower categories of objects along with endorsing sometimes stricter and sometimes laxer requirements on membership to those categories. The literature also shows that cognizers will vary which characteristics they consider to be most typical of the objects of a concept in different circumstances as well as which of those objects they consider to be most representative of a category as a whole. Indeed, it's difficult to find examples of cognition where context does not have at least some effect on which of a concept's contents are activated for the purposes of producing some act of cognition.

These empirical findings thus confirm what the appeal to cognitive efficiency would predict: cognition tends to proceed by relying on contextually relevant contents, representing different contents for the performance of different cognitive acts.¹⁰⁴ I take these considerations to show that cognition, as we know it, tends toward *variability* over *plurality*. Cognition depends on the representation of many different contents in many different contexts, tending toward the

by a revision to the Act that extends its application to endangered insects, but a simpler and faster solution was devised by simply applying the legal definition of fish in the United States, which already includes invertebrates (see again the definition quoted in fn. 93 above), to bees and endangered insects in general. This is an absurd decision only if biological conceptions of fish are given priority over legal conceptions across all contexts, a point I've already dismissed.

¹⁰⁴ I do not mean that context *necessitates* the activation of certain contents over others, only that it *motivates* the activation of relevant contents over irrelevant contents.

activation of only those select contents that are required or otherwise motivated for representation by the contextual particularities of particular cognitions. And so, as the contexts of cognition vary, so too will the contents represented for their performance. Cognition depends on concepts admitting of variability in the activation of their contents.

5.3.5 Cognition Requires Variable Storage of Representational Content

Variability in activation does not entail variability in storage.¹⁰⁵ A concept might admit of variation in the contents it activates for cognition, while always activating a selection from the same invariant store of contents. That is, even if it's true that FISH allows the activation of a variety of different contents across different acts of cognition, it might still never vary in the contents it stores for such activation. Perhaps it always represents the same invariant store of fish exemplars and fish characteristics, only admitting of variation in which of these it activates on any particular cognitive occasion.

As I argued in the previous Chapter, I don't think this is a plausible possibility; I think cognition depends on variability not only in a concept's activated contents, but also variability in its stored contents. And the reason here is the following: I think cognition requires that the stored contents of concepts be subject to *revision*. I think we should understand our concepts as capable of evolving over time as we gain experience and expertise with a concept's referents and their properties, updating the categories and the characteristics stored in a concept over time. I also think that cognition goes well to the extent that our concepts actually undergo such updates. Indeed, this

¹⁰⁵ But, since I've assumed activation is always activation of stored content, variability in storage does entail variability in activation. As stored contents vary, which contents are eligible for activation must vary too. To this extent then, my argument for variable storage bolsters my argument for variable activation.

is a most natural way of articulating what it is to *learn* a concept, and to achieve increasing *competence* and *sophistication* in its use across various contexts of cognition. But concepts can be updated only if their stored contents can vary. And so, if cognition depends on cognizers updating the contents of their concepts in light of their growing body of experiences, then concepts must admit of variation in their stored contents in order to enable cognition.

We do not, at the moment of concept acquisition, already acquire all those contents that we will eventually come to store with that concept over time. We do not immediately acquire a representation of each referent and each characteristic that will we eventually have at our disposal for the sake of performing various future acts of cognition, if only because those contents are not yet known to us. The possibility of those future acts of cognition thus depends on the acquisition of further contents, and perhaps the removal of those contents that would have made those future acts impossible.

Suppose your initial concept FISH does not include, for example, *lampreys*, and suppose you initially thought of fish as simple creatures, incapable of such neural sophistications as *pain* and *basic arithmetic*. Your eventual ability to think of *lampreys* by means of FISH then depends on the concept being updated to include such jawless and scaleless organisms. So too does your eventual ability to project the capacities to *feel pain* and *perform basic arithmetic* to lampreys, as well as other fish. And these projections are to be explained by appeal to FISH only if the concept has been updated to include such characteristics. That is, to the extent that these contents are not part of the initial contents stored by your concept FISH, they must be added to the concept if it is to be invoked to explain cognitions that depend on these contents.

In fact, even though variable activation does not entail variable storage, I do think variations in storage can be explained by variations in activation in the following sense. I've

already argued that concepts tend toward activating only contextually relevant contents. And now I've suggested that concepts can be revised over time, admitting of updates according to ongoing experience. Suppose now that concept revision proceeds via a form of *reinforcement learning*, according to which contents that are usefully activated more often are more likely to be maintained in storage than contents that are usefully activated less often, either because they aren't activated very often or because they aren't of much use when they are activated.¹⁰⁶ If this is right, then it follows that the contents stored by a concept over time will change as we move across different contexts, which prime the activation of different contents.¹⁰⁷ And, insofar as our resulting knowledge bases can differ, both between and within different cognizers over time, as a resulting of moving through different contexts, which contents are stored in our concepts will vary both inter-personally and intra-personally over time. In this sense then, variations in activation can be understood as reliably producing variations in storage. And so, to the extent that you have been convinced that concepts admit of variable activation, you should expect variable storage.

5.3.6 Interim Objection 3: Why Not Separate Storage?

I've argued that our concepts enable acts of cognition in virtue of admitting variation in their representational contents across different acts of cognition. I've extended this claim to apply to the contents stored by a concept, arguing that concepts admit of variable storage, storing different contents between different cognizers and the same cognizers over time. But a tempting

¹⁰⁶ Some empirical support for such a process is discussed in (Barsalou L. , 1982).

¹⁰⁷ How might contents be added to a concept in the first place? I think one part of the story will involve an appeal to *ad hoc* contents, which are active in virtue of being derived rather than being activated from storage. If a derived content is useful enough often enough, it might be "upgraded" to storage. The reverse might be true for useless contents. Once forming part of a concept's store of contents, sufficiently useless contents might be "downgraded", becoming unavailable for activation, only derivation.

alternative account remains possible. It remains open to my opponent to hold that our concepts are simply *replaced* with new ones whenever we learn some new information that motivates an update in the stored contents of a concept; it remains possible that each time we learn something new about fish, something which motivates an update to the concept of fish, we create a new concept, replacing FISH_n with FISH_{n+1}. A more moderate option is that we sometimes *bifurcate* our concepts upon learning that the concept can apply to a variety of different categories, i.e. bifurcating some earlier concept KID-FISH, or ADULT-NOVICE-FISH into several distinct concepts BIOLOGICAL-FISH, CULINARY-FISH, LEGAL-FISH, etc. upon learning of the different conceptions employed in different domains. These may then be further bifurcated once we learn about different conceptions within ichthyology, across different culinary cultures, and in different legal jurisdictions, as well as different historical periods. That is, one way of interpreting much of what I've discussed so far is as concerning differences that exist *between* the contents stored by *different* concepts of fish rather than any differences within any one concept FISH.

As I already argued in the previous Chapter, however, to the extent that these bifurcations involve a commitment to *Invariantism*, the move is subject to a very troubling dilemma: the resulting account of concepts either fails to offer a full accounting of the cognitive import of concepts or it risks empirical invalidity. And if it is accompanied by a commitment to *Contextualism*, the idea that concepts can persist across contexts and be shared across cognizers disappears. But even absent a commitment to these objectionable forms of *Concept Stability*, I think the move to bifurcate remains problematic, for the following reasons.

Once again, I think we can reject the counterproposal by appeal to cognitive efficiency, bolstered by some empirical findings. I think it makes better sense, from the perspective of cognitive efficiency, that we store information about closely interrelated categories of objects

132

together, in order to more quickly allow access to interrelated information that might be of relevance across different domains and that we update this information as we continue to learn more about the objects that occupy those domains.

In this respect it is worth noting the extent to which different fish-involving domains interact: the culinary domain and the legal domain might emphasize different characteristic properties of fish, but they are crucially interrelated insofar as much of the regulation and restriction on fishing, fisheries, and other fish-impacting activity in and around the water concerns the routes by which fish are taken, transported, and transformed for the sake of human consumption. And the impact these human activities have on fish and their habitats is both causally influenced by and causally impacts their biology: their life cycles, their habitats, and their everdecreasing capacity to survive and reproduce. The culinary and legal domains are thus crucially interrelated with the biological domain. Historical domains, meanwhile, are interested in precisely how the relationships between these domains have developed and changed over historical time. To the extent that cognizers might cross over these disciplinary boundaries in their interdisciplinary cognitions, a single concept FISH, admitting of several different sorts of fishy contents, is rather useful. For it allows us to keep thinking of fish as constituting a single subject matter, but exactly which fish and which of their characteristics we think and talk about on any given occasion, can vary and change over time. On my view this is precisely because which particular contents represented by the concept FISH can vary and change over time.

On the counterproposal, such variations and changes are to be explained by an appeal to several interrelated but *distinct* concepts of fish. In some ways this dispute might seem mostly verbal; it looks to be a dispute over which bundles of contents deserve the name "concept". And so I don't think our intuitions about any particular example will be able to settle the matter once

and for all. But I do think there is at least one empirical literature that supports my prediction that interrelated information about overlapping categories tends to be stored together under a single representation rather than distinct representations, and thus supports my claim that varying contents can belong to one and the same concept.

Philosophers don't usually distinguish between the phenomena of *polysemy* and *mere homonymy*, tending to reduce the former to the latter.¹⁰⁸ But empirically, there is a robust and measurable difference.¹⁰⁹ Mere homonyms are cases of ambiguity in which a word admits of two or more *unrelated* senses, e.g. "ray" is ambiguous between light rays and fish rays, objects that have little to do with one another. And empirically, we have evidence that these senses *compete* with one another; cognizers will move to resolve the ambiguity in order to parse such a homonymous term and once disambiguated, the rejected sense quickly decays, being discarded from working memory. Polysemes, by contrast, are ambiguities of *related* senses, e.g. "fish" is ambiguous between the animal and its flesh. And in this case, the senses needn't compete; cognizers needn't resolve the ambiguity in order to parse the ambiguous term, and the related senses may be retained in working memory. That is, while polysemous senses *co-activate*, homonymous senses *de-activate*.

Many theorists of polysemy agree that this difference is best interpreted as a difference between single representations and multiple representations, with the related senses of polysemes being stored in the same representation while the unrelated senses of homonyms being stored

¹⁰⁸ Consider, for example, the entry on "ambiguity" in the *Stanford Encyclopedia of Philosophy*, (Sennet, 2021), which mentions polysemy only to collapse its distinction with mere homonymy. (Fodor J., 1998, p. 53), puts the point even more strongly: "there is no such thing as polysemy".

¹⁰⁹ See, e.g. the surveys in (Vicente A., Polysemy and Word Meaning: An Account of Lexical Meaning for Different Kinds of Content Words, 2018) and (Vicente & Falkum, 2017). See also (Quilty-Dunn, 2021) who, like me, leans on the polysemy literature to defend an account of concepts according to which concepts admit of variable reference.
separately in distinct representations. And some, like me, interpret these findings as exhibiting the existence of variable concepts:

In short, psycholinguistic and neurolinguistic evidence suggests that polysemous words, unlike homonyms, are easier and quicker to retrieve and that ambiguity resolution for polysemes involves modulation of a common meaning rather than selection among competing concepts. Thus polysemy, unlike homynymy [sic], involves *a single concept with multiple available denotations*. (Quilty-Dunn, 2021, pp. 10, my emphasis)

My arguments here vindicate this interpretation of the empirical data, for my view holds that the central representational units of thought, i.e. concepts, store a variety of interrelated contents for a variety of interrelated cognitive uses, with no particular contents serving to enable all of a concept's uses across all the acts of cognition that it serves to produce. Indeed, the idea that concepts typically admit of variation in their contents offers a natural explanation of why languages so regularly admit of polysemy.¹¹⁰ Moreover, by tying my account to this body of research, I also render my interpretations of concepts like FISH testable. If FISH really does admit of variable reference and variable characterization, "fish" should pass the standard empirical tests for polysemy. Should it fail these tests, I will concede that what I've called "conceptions" of fish actually constitute distinct concepts of fish. But, given the levels of relation that exist between these conceptions, my bet is on the variability of FISH, and all the many concepts like it.

¹¹⁰ (Vicente & Falkum, 2017) speculate that "virtually every word is polysemous to some extent."

5.4 Concept Variability is a Feature, Not a Bug

On my view, concepts admit of variable characterizations of variable referents. I've called this view *Concept Variability* and I've spent the bulk of this Chapter arguing that it follows from the assumption that concepts enable cognition, i.e. from the claim that concepts explain how cognizers perform particular acts of cognition. I've leaned especially hard on the claim that it is the *contents* of concepts that enable the acts of cognition and thus that concepts admit of variations in their content to the extent that different acts of cognition depend on different contents. I also leaned especially heavily on the idea that cognition is *efficient* to the extent that it depends on variable concepts. The implication then is that variable concepts are a mark of cognition going well and thus that variability in a concept's content is a *feature* of our cognitive architecture rather than, as traditional philosophical orthodoxy would suggest, a *bug*. And so, even though I've focused on only one particular concept here, the claim of *Concept Variability* quickly generalizes across concepts of all sorts. If concepts really are the basic units of thought, we should be surprised to find any necessarily invariant.

6.0 Coda: Identity and Communication Without Stability

The adoption of *Concept Variability* opens up several exciting questions, some much harder to answer than others. I close with some speculative remarks about the question of a variable concept's *identity*, as well as the risk it might seem to pose to *communication*.

6.1 On A Variable Concept's Identity

I've denied that concepts can be individuated by any necessarily invariant contents, for on my view concepts represent only variable contents. But this does not mean that concepts cannot be individuated. Indeed, the idea that concepts are *dynamic* suggests a *historical* metaphysics, according to which a concept is to be identified by the history of its varied uses and the various contexts that shaped it along the way.

One such historical account has recently been offered by (Sainsbury & Tye, 2012). On their view, both the epistemic contents and the intentional contents of a concept can change over time. As they explicate their account, in terms of MEAT and MANSION:

The concept MEAT used in the fifteenth century had as its referent anything edible; in our terms, its content was the property of being edible. The concept MEAT we use now has as its referent only flesh; in our terms, its content is the property of being flesh. One option is to say the earlier concept is the same as the current concept, but that its content has changed. Another option is to say that a new concept, expressed by a word spelled and pronounced the same way, was introduced at some point, and each of the two concepts have retained their original and distinct contents. We prefer the first story. It is a case of gradual drift, with no event that seems a good candidate for the introduction of a new concept. (Sainsbury & Tye, 2012, p. 46)

A possibly more striking case: seventy years ago a two-story five-bedroom house might have qualified as an instance of the concept MANSION. Thanks to gradual upsizing, this is no longer so. But there seems no case for saying the current concept MANSION is distinct from the earlier one. (Sainsbury & Tye, 2012, pp. 46, fn. 9)

Concepts thus cannot be identified by their contents, for these can always change over time. Instead, on their view, concepts are to be identified by their *originating use*. According to this view then, which they have dubbed *Originalism*, each concept is used either for the very first time, in which case the concept has been *created*, or it is used in deference to one or more previous uses, in which case it has *persisted*, forming part of an ancestral lineage of ongoing use. As they put it:

According to originalism, every concept has exactly one originating use, and every originating use of a concept is an originating use of just one concept. The following is a necessary and sufficient condition for concepts to be the same:

(O) Necessarily: concept C1 = concept C2 iff the originating use of C1 = the originating use of C2.

Following an originating use, there are (zero or more) "descendant" uses: uses that count as uses of that very same concept thanks to standing in the ancestral of a deference-involving relation R to the originating use. We first came into contact with the concept QUARK at a lecture by Gell-Mann. When we first used it, we were trying to use it as he had used it in the lecture. We have maintained those efforts. Maybe we have now forgotten that early learning experience, but in trying today to use the concept as we used it yesterday we are maintaining a link of R-dependence with our first use; and so, indirectly, with Gell-Mann's originating use. (Sainsbury & Tye, 2012, p. 44)

One very important part of the *Originalist* account is that historical deference can be *partial*. We can depart more or less drastically from (our interpretation) of previous usages, while still intending to continue the ancestral lineage from which our descendant usages take their cue. And this means that we can contribute to the continuation of a conceptual lineage, without representing all the same contents represented in previous uses. It is possible for us to deliberately tinker with previously represented contents, and we can mistakenly depart from them. All that *Originalism* requires is a deferential intention to re-use a particular pre-existing concept. To this extent, I am in agreement with Sainsbury and Tye, and happily adopt their metaphysics for concept identity without objection.

Nevertheless, I have one serious reservation. Another important part of *Originalism* is that concepts with shared contents, as well as shared linguistic expressions, can nevertheless be distinct just to the extent that they have distinct originating uses. It is fully possible, according to this view, that my use of "meat" for meat, and your use of "meat" for meat, as well as my use of "vleis" and your use of "meat", express distinct concepts of meat to the extent that we (perhaps unknowingly) defer our usage to distinct historical legacies, regardless of how intertwined these legacies have become. Indeed, the legacies may be so intertwined, that each of us defers (again, perhaps unknowingly) to both (or more) lineages with distinct origins. It is also fully possible for us to create our own alternatives to existing concepts, insofar as we start thinking about some subject matter, like meat, without deferring to existing meat-thoughts, perhaps in virtue of interacting with some meat before becoming aware of how other people think about the same stuff.

In both cases, Sainsbury and Tye predict that these *distinct* concepts will persist despite appearing to be shared, or that one will be *supplanted* by the other, as one gains greater deferential

dominance over all others. Here they discuss the cases of CAT, QUARK/ACES, and INTEGRATION:

According to originalism, the concepts infants form on their own are typically supplanted by public concepts when they become full members of their surrounding linguistic community. An infant might originate a concept for cats. But as the child becomes a member of his conceptual community, that concept will be supplanted by one or more public concepts, most likely the concept CAT. (Sainsbury & Tye, 2012, p. 43)

[...] some people take Gell-Mann and Zweig to have created the concept QUARK independently, though it was Gell-Mann's word for the concept that became prevalent (Zweig's was "aces"). Likewise, there is a view according to which the concept INTEGRATION was introduced independently by Newton and by Leibniz. According to originalism, these ways of describing what happened must be incorrect, since concepts are individuated by their origin. [...] We suspect, however, that it is more historically accurate to say that deference to Gell-Mann trumps deference to Zweig. If so, the right thing to say is that there is a single concept QUARK introduced by Gell-Mann. Zweig introduced a distinct concept, the concept ACES. His conception of aces was very similar to Gell-Mann's conception of quarks. But Gell-Mann interacted with more people, was more widely cited, and won the Nobel prize, so it was his concept QUARK, rather than Zweig's concept ACES, that become prevalent. For all we know, a story with this structure might be the right one for Newton, Leibniz and the concept INTEGRATION. (Sainsbury & Tye, 2012, pp. 67-68)

Sainsbury and Tye thus rule out the possibility of conceptual *fusion*.¹¹¹ Indeed, the fusion of two conceptual lineages is treated as *metaphysically* impossible—note again that "Necessarily"

¹¹¹ They do allow conceptual *fission*: see discussion in (Sainsbury & Tye, 2012, pp. 66-67). In this case, one existing concept gives way to two newly created concepts, each with their own originating use. The result is either three distinct concepts, if the old one persists alongside the two new concepts, or the result is two new concepts, if the old one ceases to be the subject of deference in future usage.

in their identity conditions for concept identity. And yet their defense of this impossibility reduces to what looks to me to be a rather ad hoc interpretation of a history of empirical events: infants sometimes begin by employing their own concepts of familiar subject matters but always end up abandoning their own original concepts and forming the intention to defer to the concepts of a wider conceptual community, a community that is presupposed to defer uniformly to a single historical lineage of cat-thought, much as the wider conceptual community eventually came to uniformly defer to Gell-Mann's QUARK over Zweig's ACES, as well as to someone's INTEGRATION, even though I for one cannot say whether we've decided to defer to the lineage that started with Newton or to the one that started with Leibniz. I certainly don't see any reason, other than a pre-existing commitment to *Originalism*, why we cannot say that at least sometimes we have started to defer to a *plurality* of intertwined lineages, or even to a *merged* lineage of use.

Originalism is already committed to the idea that we defer to a *community* of cognizers, i.e. to the usage observed across a variety of different users both past and present, many of whom needn't represent the same contents through the use of their shared concept, at least across all contexts of use. If the deference can tolerate a plurality of uses, users, contents, and contexts, why can't it also tolerate a plurality of perhaps unknown historical origins? If we as a community can defer sometimes more and sometimes less faithfully to one another's cat-thought, quark-thought, integration-thought, meat-thought, and so on, in our own thoughts, why can't that be enough for us to share concepts like CAT, QUARK, and MEAT? What more does a shared origin offer?

My own view would involve generalizing away from the strict originalism of Sainsbury and Tye, insofar as I hold that concepts can be shared across different disciplines, cultures, and other domains which needn't defer to a single historical legacy of a concept with a single historical origin. I think that different variants of the same concepts can originate separately, with their initially separate historical lineages being available for a "fusion" or "merge" upon inter-personal, inter-disciplinary, or inter-cultural contact.

On such an account, it would be possible that both Newton and Leibniz devised a version of INTEGRATION, variants that soon came to be merged into a single shared concept to which we all now defer. It would remain possible that Zweig's ACES came to be rejected in favor of Gell-Mann's QUARK, but it is also possible that it was incorporated into the lineage of usage that defers primarily to Gell-Mann's QUARK. It would also be possible for infants to gradually adjust their own CAT to better conform to the community of CAT-users that help them learn about cats, intermingling their own lineage of use with the uses of their wider community. And it would be possible for separate communities across the world to have developed their own variants of concepts, including CAT and MEAT, as well as BIRD, FRUIT, FISH, and more, many of which have since been intertwined and even merged, at least partially, as a result of our ever-increasing inter-disciplinary and cross-cultural contact and collaboration.

6.2 On Using a Variable Concept in Communication

I'd like to close this discussion with a closer look at the idea that variable concepts can enable successful communication. I think it's rather common to assume that successful communication requires not only that we employ the same concepts, but that those concepts represent the same contents. Otherwise, we are simply "talking past one another".

Recall again Barsalou's articulation of the idea:

[...] intuitions of stability are often compelling, and it is sometimes hard to imagine how we could communicate or perform other intelligent behaviors without *stable* knowledge structures [i.e. concepts]. (Barsalou L., 1989, p. 76; my emphasis)

As Hampton puts it, in terms of FISH:

When we communicate with others, the very definition of successful communication is clearly going to depend crucially on our ability to know when we are talking about the same thing. Otherwise we might not be speaking the same language, and disagreement and agreement would be meaningless. If by *fish* I mean what biologists mean by fish, while Ishmael means something like "fish-shaped creatures of the deep," unless we are aware that we have these different concepts, we will be open to all kinds of misunderstanding. (Hampton J., 2015, p. 666)

In line with the intuitions identified by Barsalou, Hampton assumes here not only that the representation of different contents implies distinctness of concepts, but he also assumes that such divergence in content necessarily precludes successful communication. At the very least, Hampton holds that we must represent the same *intentional* contents: we must mean the same *fish* when using the concept FISH. On my view, this requirement is too strict. *Concept Variability* cannot imply that we cannot communicate.

In my view, communication does not require that we represent the same contents while using a shared concept. I think we can get away with a mere *overlap*. In particular, I think communication requires only that we converge on contextually relevant contents, whether those are epistemic, intentional, or both. In order to get this idea on the table, consider this lovely example offered by Bezuidenhout: [...] my son wants to read the daily comics and is searching for the newspaper. I tell him: 'I left the newspaper on the dining room table.' In fact what is on the table is this morning's paper minus the business section, which his father has taken to his office. But I know my son won't even notice that the business section is missing, since all he will look for is the section that has the comics, and that section is amongst those on the table. So for current purposes the newspaper is on the table, despite the fact that a complete copy is not on the table. (Bezuidenhout, 2002, p. 108)

Her son has inquired about the location of the daily newspaper, i.e. the *entire* newspaper. Bezuidenhout responds by telling him that the newspaper is on the table, glossing over the fact that it is only a *part* of the newspaper that is on the table. And yet, despite this referential divergence in what is meant by "the newspaper" between mother and son, no risk of misunderstanding looms, since the mutually known purpose of the exchange is to acquire information about the location of the comics section, which is on the table.

In order to emphasize that the purpose of the exchange can affect whether a referential divergence results in miscommunication, Bezuidenhout offers a second scenario:

Suppose that unbeknownst to me the reason my son was searching for the newspaper is that he wanted to check the financial section to see the stock market listings for a school project he is working on. When he finds the business section is missing he might accuse me of misleading him when I said the newspaper was on the table, since the entire newspaper was not on the table. However, this is just a normal case of miscommunication. I thought (mistakenly) that my son was looking for the comic section of the paper. I was unaware of a crucial piece of information that would have explained his quest, namely that he was working on his school project. This lack of mutual knowledge is what explains the different understandings that we have of the crucial term 'newspaper'. (Bezuidenhout, 2002, pp. 108-109)

In this case, it is contextually relevant that a part of the newspaper is missing, since this is precisely the part of the newspaper that the son seeks. In this way, the two cases illustrate that referential divergences can indeed result in miscommunication, but only when those referential divergences are contextually relevant. As long as referential divergence is contextually irrelevant, and contextually relevant referential convergence is secured, communication can succeed.

I think we can straightforwardly extend Bezuidenhout's case to variable concepts. Variable concepts pose no threat to successful communication, as long as we converge on contextually relevant epistemic and intentional contents. If you mean what biologists mean by "fish" while I use it to mean something like "fish-shaped creatures of the deep", we may still proceed successfully in a conversation about fish, as long as our conversation doesn't depend for its success on matters concerning *aquatic mammals*. In much the same way, an ichthyologist can order "the fish" for their dinner without showing any surprise when the waiter brings them a plate of *scallops*. We needn't represent the same contents by means of the shared concept FISH, as long as the contents we do represent on the occasion of our exchange converges on those contents that are relevant to the purposes of our exchange, i.e. as long as any diverging representational contents are contextually irrelevant.

On this view then, it is true that communication can fail even when we employ the same concepts, and it is true that communication can fail because we use the same concept to represent different contents, but it is not true that communication can fail just in virtue of using the same concept to represent different contents. The mere fact that concepts can vary poses no risk to communication. The risk comes from *contextually misaligned* contents, resulting from being ill-informed of your interlocutor's communicative intent or the purpose of your exchange, and thus

of some aspects of the context in which the exchange is taking place. But there's no controversy in that.

Bibliography

- Abad-Santos, A. (2014, December 17). *New York law thinks a burrito is a sandwich*. Retrieved July 2020, from Vox: https://www.vox.com/2014/7/22/5877325/new-york-law-sandwich-burrito-tax
- Abbott, J., & Kemp, C. (2020). Birds and Words: Exploring Environmental Influences on Folk Categorization. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 42, 744-749.
- Ahn, W., & Kim, N. (2001). The Causal Status Effect in Categorization: An Overview. *The Psychology of Learning and Motivation: Advances in Research and Theory*, 40, 23-65.
- Ahn, W., Kalish, C., Gelman, S., Medin, D., Luhmann, C., Atran, S., . . . Shafto, P. (2001). Why Essences are Essential in the Psychology of Concepts. *Cognition*, 82, 59-69.
- Ahn, W., Marsh, J., Luhmann, C., & Lee, K. (2002). Effect of Theory-Based Feature Correlations on Typicality Judgments. *Memory & Cognition, 30*, 107-118.
- Allott, N., & Textor, M. (2012). Lexical Pragmatic Adjustment and the Nature of Ad Hoc Concepts. *International Review of Pragmatics*, 4(2), 185–208.
- Ameel, E., & Storms, G. (2006). From Prototypes to Caricatures: Geometrical Models for Concept Typicality. *Journal of Memory and Language*, *55*(3), 402-421.
- Amin, T., & Levrini, O. (Eds.). (2018). Converging Perspectives on Conceptual Change. Routledge.
- Anderson, J., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin & Review*, 8, 629-647.
- Anderson, R., & Ortony, A. (1975). On Putting Apples in Bottles: A Problem of Polysemy. *Cogntive Psychology*, 7, 167-180.
- Anderson, R., & Shifrin, Z. (1980). The Meaning of Words in Context. In R. Spiro, B. Bruce, & W. Brewer (Eds), *Theoretical Issues in Reading Comprehension: Perspectives from Cognitive Psychology, Linguistics, Artificial Intelligence and Education* (pp. 331-348). Taylor and Francis.
- Anderson, R., Pichert, J., Goetz, E., Schallert, D., Stevens, K., & Trollip, S. (1976). Instantiation of General Terms. *Journal of Verbal Learning and Verbal Behavior*, *15*, 667-679.
- Anglin, J. (1977). Word, Object, and Conceptual Development. WW Norton.

- Armstrong, S., Gleitman, L., & Gleitman, H. (1983). What Some Concepts Might Not Be. *Cognition*, 13, 263-308.
- Ashcraft, M. (1978a). Property Norms for Typical and Atypical Items from 17 Categories: A Description and Discussion. *Memory & Cognition*, 6(3), 227-232.
- Ashcraft, M. (1978b). Property Dominance and Typicality Effects in Property Statement Verification. *Journal of Verbal Learning and Verbal Behavior*, 17(2), 155-164.
- Atran, S. (1999). Itzaj Maya Folk-Biological Taxonomy. In D. Medin, & S. Atran (Eds), *Folkbiology*. MIT Press.
- Baddeley, A., Eysenck, M., & Anderson, M. (2009). Memory. Psychology Press.
- Bailenson, J., Shum, M., Atran, S., Medin, D., & Coley, J. (2002). A Bird's Eye View: Biological Categorization and Reasoning Within and Across Cultures. *Cognition*, 84, 1-53.
- Balcombe, J. (2017). What a Fish Knows: The Inner Lives of Our Undeerwater Cousins. Scientific American / Farrar, Straus and Giroux.
- Barclay, J., Bransford, J., Franks, J., McCarrell, N., & Nitsch, K. (1974). Comprehension and Semantic Flexibility. *Journal of Verbal Learning and Verbal Behavior*, *13*, 471-481.
- Barsalou, L. (1982). Context-Independent and Context-Dependent Information in Concepts. *Memory & Cognition, 10,* 82-93.
- Barsalou, L. (1983). Ad hoc categories. *Memory & Cognition*, 11(3), 211-227.
- Barsalou, L. (1985). Ideals, Central Tendency, and Frequency of Instantiation as Determinants of Graded Structure in Categories. *Journal of Experimental Psychology*, *11*(4), 629-654.
- Barsalou, L. (1987). The Instability of Graded Structure: Implications for the Nature of Concepts. In U. Neisser (Ed.), Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization (pp. 101-140). Cambridge: Cambridge University Press.
- Barsalou, L. (1989). Intraconcept Similarity and its Implications for Interconcept Similarity. In S. Vosniadou, & A. Ortony, *Similarity and Analogical Reasoning* (pp. 76-121). Cambridge: Cambridge University Press.
- Barsalou, L. (1993). Flexibility, Structure, and Linguistic Vagary in Concepts: Manifestations of a Compositional System of Perceptual Symbols. In A. Collins, M. Gathercole, M. Conway, & P. Morris (Eds), *Theories of Memory* (pp. 29-101). Lawrence Erlbaum Associates.
- Barsalou, L. (2003). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 358, 1177–1187.
- Barsalou, L. (2005). Abstraction as Dynamic Interpretation in Perceptual Symbol Systems. In L. Gershkoff-Stowe, & D. Rakison (Eds), *Building object categories* (pp. 389-431). Erlbaum.

- Barsalou, L., & Sewell, D. (1984). Constructing Representations of Categories From Different Points of View. Manuscript.
- Barsalou, L., Wilson, C., & Hasenkamp, W. (2010). On the Vices of Nominalization and the Virtues of Contextualizing. In B. Mesquita, L. Feldman Barrett, & E. Smith, *The Mind in Context* (pp. 334-360). Guilford Press.
- Battig, W., & Montague, W. (1969). Category Norms for Verbal Items in 56 Categories: A Replication and Extension of the Connecticut Category Norms. *Journal of Experimental Psychology*, 80(3), 1-46.
- Beilock, S., Lyons, I., & Mattarella-Micke, A. (2008). Sports experience changes the neural processing of action language. *Psychological and Cognitive Sciences*, 105(36), 13269-13273.
- Bellazza, F. (1984b). Reliability of Retrieval from Semantic Memory: Noun Meanings. *Bulletin* of the Psychonomic Society, 22(5), 377-380.
- Bellezza, F. (1984a). Reliability of Retrieval from Semantic Memory: Common Categories. *Bulletin of the Psychonomic Society*, 22(4), 324-326.
- Belohlavek, R., & Klir, G. (2011). Concepts and Fuzzy Logic. Cambridge: MIT Press.
- Bezuidenhout, A. (2002). Truth-Conditional Pragmatics. Philosophical Perspectives, 16, 105-134.
- Bisanz, G., LaPorte, R., Vesonder, G., & Voss, J. (1978). On the Representation of Prose: New Dimensions. *Journal of Verbal Learning & Verbal Behavior*, 17(3), 337-357.
- Block, N. (1995). How many concepts of consciousness? *Behavioral and Brain Sciences*, 18(2), 272-287.
- Borkenau, P. (1990). Traits as Ideal-Based and Goal-Derived Social Categories. *Journal of Personality and Social Psychology*, 58(3), 381-396.
- Boroditsky, L. (2001). Does Language Shape Thought?: Mandarin and English Speakers' Conceptions of Time. *Cognitive Psychology*, 43(1), 1-22.
- Braisby, N. (1993). Stable Concepts and Context-Sensitive Classification. *The Irish Journal of Psychology*, *14*(3), 426-441.
- Braisby, N., & Franks, B. (2000). Perspectives and Categorization. Manuscript.
- Braisby, N., Franks, B., & Hampton, J. (1996). Essentialism, Word Use, and Concepts. *Cognition*, 59(3), 247-274.
- Braisby, N., Franks, B., & Harris, J. (1997). Classification and Concepts: Fuzzy or Perspectival? In D. Dubois (Ed.), *Catégorisation et Cognition: de la Perception au Discours* (pp. 163-189). Kimé.

- Brigandt, I. (2010). The epistemic goal of a concept: accounting for the rationality of semantic change and variation. *Synthese*, 177, 19-40.
- Brooks, L. (1987). Decentralized Control of Categorization: The Role of Prior Processing Episodes. In U. Neisser (Ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization* (pp. 141-174). Cambridge: Cambridge University Press.
- Brown, C. (2013). Hand and Arm. In M. Dryer, & M. Haspelmath (Eds), *WALS Online (v2020.3)*. Zenodo.
- Burnett, R., Medin, D., Ross, N., & Blok, S. (2005). Ideal is Typical. *Canadian Journal of Experimental Psychology*, 59(1), 3-10.
- Camp, E. (2015). Logical Concepts and Associative Characterizations. In E. Margolis, & S. Laurence, *The Conceptual Mind* (pp. 591-621). Cambridge, Massachusetts; London, England: The MIT Press.
- Carey, S. (2009). The Origin of Concepts. Oxford : Oxford University Press.
- Carnap, R. (1950). Logical Foundations of Probability. The Unversity of Chicago Press.
- Carston. (2002). Thoughts and Utterances: The Pragmatics of Explicit Communication. Blackwell.
- Casasanto, D., & Lupyan, G. (2015). All Concepts Are Ad Hoc Concepts. In E. Margolis, & S. Laurence, *The Coneptual Mind* (pp. 543-566). Cambridge, Massachusetts; London, England: The MIT Press.
- Coleman, L., & Kay, P. (1981). Prototype Semantics: The English Word Lie. Language. *Language*, 57(1), 26-44.
- Currie, A., & Killin, A. (2016). Musical pluralism and the science of music. *European Journal for Philosophy of Science*, 6(1), 9-30.
- Davis, T., & Love, B. (2010). Memory for Category Information Is Idealized Through Contrast With Competing Options. *Psychological Science*, *21*(2), 234-242.
- Del Pinal, G., & Reuter, K. (2015). "Jack is a True Scientist": On the Content of Dual Character Concepts. *Proceedings of the Annual Meeting og the Cognitive Society*, *37*, 554-559.
- Del Pinal, G., & Reuter, K. (2017). Dual Character Concepts in Social Cognition: Commitments and the Normative Dimension of Conceptual Representation. *Cognitive Science*, *41*, 477-501.
- Dieciuc, M., & Folstein, J. (2018). Typicality: Stable Structures and Flexible Functions. *Psychonomic Bulletin and Review*, 26, 491-505.

- Douven, I., & Gärdenfors, P. (2019). What are Natural Concepts? A Design Perspective. *Mind and Language*, *35*(3), 313-334.
- Dupré, J. (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Harvard University Press.
- Dupré, J. (1999). Are Whales Fish? In D. Medin, & S. Atran (Eds), *Folkbiology* (pp. 461–746). MIT Press.
- Enfield, N., Majid, A., & van Staten, M. (2006). Cross-linguistic categorisation of the body: Introduction. *Language Sciences*, 28, 137–147.
- Ereshefsky, M., & Reydon, T. A. (2015). Scientific kinds. *Philosophical Studies*, 172(4), 969–986.
- Erickson, M., & Kruschke, J. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology*, 127(2), 107-140.
- Esau, K. (1999). Anatomy of Seed Plants, 2nd Edition. Wiley.
- Fee, A., Segabache, R., & Tobolski, E. (1985). Hardness Testing. In J. Newby (Ed.), *Metals Handbook, Ninth Edition*. American Society for Metals.
- Fodor, J. (1975). The Language of Thought. New York, New York: Thomas Y. Crowell.
- Fodor, J. (1998). Concepts. Oxford: Oxford University Press.
- Fodor, J., Garrett, M., Walker, E., & Parkes, C. (1980). Against Definitions. *Cognition*, 8(3), 263-367.
- Foster-Hanson, E., & Rhodes, M. (2019). Is the Most Representative Skunk the Average or the Stinkiest? Developmental Changes in Representations of Biological Categories. *Cognitive Psychology*, 110, 1-15.
- Franks, B. (1995). Sense Generation: A "Quasi-Classical" Approach to Concepts and Concept Combination. *Cognitive Science*, 441-505.
- Franks, B., & Braisby, N. (1990). Sense Generation, or How to Make a Mental Lexicon Flexible. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 12, 438-446.
- Garnham, A. (1979). Instantiation of Verbs. *The Quarterly Journal of Experimental Psychology*, 31(2), 207-214.
- Gelman, S. (2003). The Essential Child. Oxford: Oxford University Press.
- Gelman, S. (2004). Psychological Essentialism in Children. Trends in Cognitive Science, 8(9), 404-409.
- Gelman, S., & Markman, E. (1986). Categories and Induction in Young Children. *Cognition*, 23(3), 183-209.

- Gelman, S., & Markman, E. (1987). Young Children's Inductions from Natural Kinds: The Role of Categories and Appearances. *Child Development*, 58(6), 1532-1541.
- Glass, A., & Meany, P. (1978). Evidence for Two Kinds of Low-Typical Instances in a Categorization Task. *Memory & Cognition*, 6(6), 622-628.
- Goldstone, R. (1996). Isolated and Interrelated Concepts. *Memory & Cognition, 24*, 608-628.
- Gonnerman, C., & Weinberg, J. (2010). Two uneliminated uses for "concepts": Hybrids and guides for inquiry. *Behavioral and Brain Sciences*, *33*(2-3), 211-212.
- Greenspan, S. (1986). Semantic Flexibility and Referential Specificity of Concrete Nouns. *Journal* of Memory and Language, 25(5), 1986.
- Halff, H., Ortony, A., & Anderson, R. (1976). A Context-Sensitive Representation of Word Meanings. *Memory & Cognition, 4*, 378-383.
- Hampton, J. (1979). Polymorphous Conceptsin Semantic Memory. *Journal of Verbal Learning* and Verbal Behavior, 18, 441-461.
- Hampton, J. (1995). Testing the Prototype Theory of Concepts. *Journal of Memory and Language*, 34, 686-708.
- Hampton, J. (2006). Typicality, Graded Membership, and Vagueness. *Cognitive Science*, *31*, 355-384.
- Hampton, J. (2015). Concepts in the Semantic Triangle. In E. Margolis, & S. Laurence (Eds), *The Conceptual Mind* (pp. 655–676). MIT Press.
- Hampton, J., & Gardiner, M. (1984). Measures of Internal Category Structure: A Correlational Analysis of Normative Data. *British Journal of Psychology*, 74(4), 419-516.
- Hampton, J., Dubois, D., & Yeh, W. (2006). Effects of Classification Context on Categorization in Natural Categories. *Memory & Cognition*, 34, 1431-1443.
- Hastings, P., Walker, H., & Galland, G. (2014). *Fishes: A Guide to their Diversity*. University of California Press.
- Heider, E. (1972). Universals in Color Naming and Memory. *Journal of Experimental Psychology*, 93(1), 10-20.
- Helfman, G., Collette, B., Facey, D., & Bowen, B. (2009). *The Diversity of Fishes: Biology, Evolution, and Ecology, 2nd edition.* Blackwell.
- Hendry, R. (2008). Two conceptions of the chemical bond. Philosophy of Science, 75, 909-920.
- Hintzman, D. (1986). "Schema Abstraction" in a Multiple-Trace Memory Mode. *Psychological Review*, 93(4), 411-428.

- Hintzman, D., & Ludlam, G. (1980). Differential Forgetting of Prototypes and Old Instances: Simulation by an Exemplar-Based Classification Model. *Memory and Cognition*, 8(4), 378-382.
- Hoenig, K., Müller, C., Herrnberger, B., Sim, E.-J., Spitzer, M., Ehret, G., & Kiefer, M. (2011). Neuroplasticity of semantic representations for musical instruments in professional musicians. *Nueroimage*, 56(3), 1714-1725.
- Hoenig, K., Sim, E.-J., Bochev, V., Herrnberger, B., & Kiefer, M. (2008). Conceptual flexibility in the human brain: dynamic recruitment of semantic maps from visual, motor, and motionrelated areas. *Journal of Cognitive Neuroscience*, 20(10), 1799-1814.
- Horstmann, G. (2002). Facial Expressions of Emotion: Does the Prototype Represent Central Tendency, Frequency of Instantiation, or an Ideal? *Emotion*, 2(3), 297-303.
- Irvine, E. (2013). Consciousness as a scientific concept: a philosophy of science perspective. Springer.
- Jackson, F. (1998). From Metaphysics to Ethics. Oxford: Claredon Press.
- Janczura, G., & Nelson, D. (1999). Concept Accessibility as the Determinant of Typicality Judgments. *The American Journal of Psychology*, 112(1), 1-19.
- Jerzak, E. (2019). Two Ways to Want? The Journal of Philosophy, 116(2), 65-98.
- Judkis, M. (2018, December 12). A hot dog is a taco. A steak is a salad. A Pop-Tart is a calzon. Let the Cube Rule Explain. Retrieved July 2020, from The Washington Post: https://www.washingtonpost.com/news/voraciously/wp/2018/12/12/a-hot-dog-is-a-tacoa-steak-is-a-salad-a-pop-tart-is-a-calzone-let-the-cube-rule-explain/
- Keefe, R. (2000). Theories of Vagueness. Cambridge: Cambridge University Press.
- Keil, F. (1989). Concepts, Kinds, and Cognitive Development. MIT Press.
- Keil, F. (1994). Explanation, Association, and the Acquisition of Word Meaning. *Lingua*, 92, 169-196.
- Kim, S., & Murphy, G. (2011). Ideals and Category Typicality. *Journal of Experimental Psychology, Learning, Memory, and Cognition, 37*(5), 1092-1112.
- Kirchin, S. (Ed.). (2013). Thick Concepts. Oxford: Oxford University Press.
- Kitcher, P. (1984). Species. Philosophy of Science, 51, 308-333.
- Knobe, J., Prasada, S., & Newman, G. (2013). Dual Character Concepts and the Normative Dimension of Conceptual Representation. *Cognition*, 127, 242-257.
- Lakatos, I. (1976). Proofs and Refutations. Cambridge University Press.

Lakoff, G. (1987). Women, Fire, and Dangerous Things. Chicago: University of Chicago Press.

- Lalumera, E. (2014). On the Explanatory Value of the Concept-Conception Distinction. *Rivista Italiana Di Filosofia Del Linguaggio*, 73–81.
- Landau, B. (1982). Will the Real Grandmother Please Stand Up? The Psychological Reality of Dual Meaning Representations. *Journal of Psycholinguistic Research*, 11(1), 47-62.
- Leslie, S.-J. (2013). Essence and Natural Kinds: When Science Meets Preschooler Intuition. In T. Gendler, & J. Hawthorne (Eds), *Oxford Studies in Epistemology, Vol. 4* (pp. 108-165). Oxford University Press.
- Leslie, S.-J. (2015). "Hillary Clinton is the Only Man in the Obama Administration": Dual Character Concepts, Generics and Gender. *Analytic Philosophy*, 56(2), 111-114.
- Levering, K., & Kurtz, K. (2006). The Influence of Learning to Distinguish Categories on Graded Structure. Proceedings of the Annual Meeting of the Cognitive Science Society, 28, 1681-1686.
- Lohr, G. (2017). Abstract concepts, compositionality, and the contextualism-invariantism debate. *Philosophical Psychology*, *30*(6), 689–710.
- Lynch, E., Coley, J., & Medin, D. (2000). Tall is Typical: Central Tendency, Ideal Dimensions, and Graded Category Structure Among Tree Experts and Novices. *Memory & Cognition*, 28(1), 41-50.
- Machery, E. (2009). Doing Without Concepts. Oxford: Oxford University Press.
- Machery, E. (2011). Concepts: A Tutorial. In R. Belohlavek, & G. Klir, *Concepts and Fuzzy Logic* (pp. 13-44). Cambridge: MIT Press.
- Machery, E. (2015). By Default: Concepts Are Accessed in a Context-Independent Manner. In E. Margolis, & S. Laurence, *The Conceptual Mind* (pp. 567-588). Cambridge, Massachusetts; London, England: The MIT Press.
- Machery, E., & Seppälä, S. (2010). Against Hybrid Theories of Concepts. Anthropology & Philosophy, 97-125.
- Majid, A., Bowerman, M., Kita, S., Haun, D., & Levinson, S. (2004). Can language restructure cognition? The case for space. *Trends in Cognitive Sciences*, 8(3), 108-114.
- Malt, B. (1994). Water is not H2O. Cognitive Psychology, 27(1), 41-70.
- Malt, B., & Smith, E. (1984). Correlated Properties in Natural Categories. *Journal of Verbal Learning and Behavior*, 23(2), 250-269.

- Margolis, E., & Laurence, S. (1999). Concepts and Cognitive Science. In E. Margolis, S. Laurence (Eds), E. Margolis, & S. Laurence (Eds.), *Concepts: Core Readings* (pp. 3-81). Bradford Books.
- Margolis, E., & Laurence, S. (2015). *The Conceptual Mind*. Cambridge, Massachusetts; London, England: The MIT Press.
- Margolis, E., & Laurence, S. (2019). *Concepts*. Retrieved from Stanford Encyclopedia of Philosophy: https://plato.stanford.edu/entries/concepts/
- Markman, A., & Ross, B. (2003). Category Use and Category Learning. *Psychological Bulletin*, 129(4), 592-613.
- Marques, T., & Wikforss, Å. (Eds.). (2020). *Shifting Concepts: The Philosophy and Psychology* of Conceptual Variability. Oxford: Oxford University Press.
- McCloskey, M., & Glucksberg, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory* & *Cognition*, 6(4), 462-472.
- Medin, D., & Atran, S. (Eds.). (1999). Folkbiology. MIT Press.
- Medin, D., & Atran, S. (2004). The Native Mind: Biological Categorization and Reasoning in Development and Across Cultures. *Psychological Review*, 111(4), 960-983.
- Medin, D., & Ortony, A. (1989). Psychological Essentialism. In S. Vosniadou, & A. Ortony, *Similarity and analogical reasoning* (pp. 179-195). Cambridge: Cambridge University Press.
- Medin, D., & Schaffer, M. (1978). Context Theory of Classification Learning. *Psychological Review*, 85(3), 207-238.
- Medin, D., Lynch, E., Coley, J., & Atran, S. (1997). Categorization and Reasoning among Tree Experts: Do All Roads Lead to Rome? *Cognitive Psychology*, *32*, 49-96.
- Mervis, C., & Pani, J. (1980). Acquisition of basic object categories. *Cognitive Psychology*, *12*(4), 496–522.
- Mervis, C., & Rosch, E. (1981). Categorization of Natural Objects. *Annual Review of Psychology*, 32, 89-115.
- Mervis, C., Caitlin, J., & Rosch, E. (1976). Relationships among goodness-of-example, category norms, and word frequency. *Bulletin of the Psychonomic Society*, 7(3), 283–284.
- Millikan, R. (2000). On Clear and Confused Ideas. Cambridge University Press.
- Millikan, R. (2005). Language: A Biological Model. Oxford: Claredon Press.
- Murphy, G. (2002). *The Big Book of Concepts*. Cambridge, Massachusetts; London, England: A Bradford Book, The MIT Press.

- Murphy, G. (2016). Is there an Exemplar Theory of Concepts? *Psychonomic Bulletin and Review*, 23, 1035-1042.
- Murphy, G., & Medin, D. (1985). The Role of Theories in Conceptual Coherence. *Psychological Review*, 92(3), 289-316.
- Neufeld, E. (2022). Psychological Essentialism and the Structure of Concepts. *Philosophical Compass*, 17(5).
- Newman, G., & Knobe, J. (2019). The Essence of Essentialism. *Mind and Language, 34*, 585–605.
- Nosofksy, R. (1991). Typicality in Logically Defined Categories: Exemplar-Similarity Versus Rule Instantiation. *Memory and Cognition*, 19(2), 131-150.
- Nosofsky, R. (1984). Choice, Similarity, and the Context Theory of Classification. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10(1), 104-114.
- Nosofsky, R. (1986). Attention, Similarity, and the Identification–Categorization Relationship. Journal of Experimental Psychology: General, 115(1), 39-57.
- Nosofsky, R. (1988). Similarity, Frequency, and Category Representations. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14(1), 54-65.
- Nosofsky, R., & Johansen, M. (2000). Exemplar-Based Accounts of "Multiple-System" Phenomena in Perceptual Categorization. *Psychonomic Bulletin and Review*, 7(3), 375-402.
- Nosofsky, R., & Zaki, S. (2002). Exemplar and Prototype Models Revisited: Response Strategies, Selective Attention, and Stimulus Generalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*(5), 924-940.
- Nosofsky, R., B., M., & Kumar, P. (2020). Contrasting Exemplar and Prototype Models in a Natural Science Category Domain. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *42*, 641-647.
- Osawa, Y., & Ellen, R. (2014). The Cultural Cognition of Taste Term Conflation. *The Senses and Society*, 9(1), 72-91.
- Osherson, D., & Smith, E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, *9*(1), 35–58.
- Osherson, D., Smith, E., Wilkie, O., López-Rousseau, A., & Shafir, E. (1990). Category-Based Induction. *Psychological Review*, 97, 185-200.
- Palmeri, T., & Nosofsky, R. (2010). Central Tendencies, Extreme Points, and Prototype Enhancement Effects in Ill-Defined Perceptual Categorization. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 54A(1), 197-235.

- Papafragou, A., Hulbert, J., & Trueswell, J. (2008). Does Language Guide Event Perception? Evidence from Eye Movements. *Cognition*, 108(1), 155-184.
- Peacocke, C. (1998). Implicit Conceptions, Understanding and Rationality. *Philosophical Issues*, 9, 43-88.
- Phillips, J., De Freitas, J., Mott, C., Gruber, J., & Knobe, J. (2017). True Happiness: The Role of Morality in the Folk Concept of Happiness. *Journal of Experimental Psychology: General*, 146(2), 165-181.
- Posner, M., & Keele, S. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3), 353-363.
- Potter, M., & Faulconer, B. (1979). Understanding Noun Phrases. *Journal of Verbal Learning and Verbal Behavior*, 18(5), 509-521.
- Prinz, J. (2002). Furnishing the Mind: Concepts and Their Perceptual Basis. Cambridge University Press.
- Quilty-Dunn, J. (2021). Polysemy and Thought: Toward a Generative Theory of Concepts. *Mind & Language*, *36*(1), 158-185.
- Quine, W. V. (1961). Two Dogmas of Empiricism. In W. V. Quine, *From a Logical Point of View,* 2nd Edition revised. New York, Hagerstown, San Fransisco, London: Harper Torchbooks.
- Regier, T., & Kay, P. (2004). Color Naming and Sunlight: Commentary on Lindsey and Brown (2002). *Psychological Science*, *15*(4), 289-290.
- Rehder, B. (2003). A Causal-Model Theory of Conceptual Representation and Categorization. Journal of Experimental Psychology: Learning, Memory, and Cognition, 29(6), 1141-1159.
- Rein, J., Goldwater, M., & Markman, A. (2010). What is Typical about the Typicality Effect in Category-Based Induction? *Memory & Cognition*, 38(3), 377-388.
- Reuter, K. (2019). Dual Character Concepts. Philosophy Compass, 14(1), e12557.
- Rey, G. (1985). Concepts and conceptions: A reply to Smith, Medin and Rips. *Cognition*, *19*, 297-303.
- Rey, G. (2010). Concepts versus Conceptions (Again). *Behavioral and Brain Sciences*, 33(2-3), 221-222.
- Rice, C. (2016). Concepts as Pluralistic Hybrids. *Philosophy and Phenomenological Research*, 92(3), 597-619.
- Rips, L. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning & Verbal Behavior*, 14(6), 665–681.

- Rips, L. (1989). Similarity, Typicality, and Categorization. In S. Vosniadou, & A. Ortony (Eds), *Simialrity and Analogical Reasoning* (pp. 21-59). Cambridge: Cambridge University Press.
- Rips, L., Shoben, E., & Smith, E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12(1), 1-20.
- Rosch, E. (1973a). On the Internal Structure of Perceptual and Semantic Categories. In T. Moore (Ed.), *Cognitive Development and Acquisition of Language* (pp. 111-144). Academic Press.
- Rosch, E. (1973b). Natural Categories. Cognitve Psychologyy, 4(3), 328-350.
- Rosch, E. (1975a). Cognitive Representations of Semantic Categories. *Journal of Experimental Psychology: General, 104*(3), 192-233.
- Rosch, E. (1975b). The Nature of Mental Codes for Color Categories. *Journal of Experimental Psychology: Human Perception and Performance, 1*(4), 303-322.
- Rosch, E. (1977). Human Categorization. In N. Warren (Ed.), Advances in Cross-Cultural Psychology (Vol. 1) (pp. 1-72). Academic Press.
- Rosch, E. (1978). Principles of Categorization. In E. L. Rosch, *Cognition and Categorization* (pp. 28-49). Elrbaum.
- Rosch, E., & Mervis, C. (1975). Family Resemblances: Studies in the Internal Structure of Categories. *Cognitive Psychology*, 7(4), 573-605.
- Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic Objects in Natural Categories. *Cogntive Psychology*, 8, 382-439.
- Rose, D., & Nichols, S. (2019). Teleological Essentialism. Cognitive Science, 43(4), e12725.
- Rose, D., & Nichols, S. (2020). Teleological Essentialism: Generalized. *Cognitive Science*, 44(3), e12818.
- Roth, E., & Shoben, E. (1983). The Effect of Context on the Structure of Categories. *Cogntive Psychology*, *15*(3), 346-378.
- Sainsbury, R., & Tye, M. (2012). Seven Puzzles of Thought and How to Solve Them. Oxford University Press.
- Savitsky, Z. (2022). Fish Can Learn Basic Arithmetic. Science.
- Sennet, A. (2021). Ambiguity. The Stanford Encyclopedia of Philosophy.
- Simoons, F. (1994). Eat Not This Flesh: Food Avoidances from Prehistory to the Present, 2nd Edition. University of Wisconsin Press.

- Smith, E., & Medin, D. (1981). *Categories and Concepts*. Cambridge, Massachusetts; London, England: Harvard University Press.
- Smith, E., Osherson, D., Rips, L., & Keane, M. (1988). Combining Prototypes: A Selective Modification Model. *Cognitive Science*, 12, 485–527.
- Smith, L., & Samuelson, L. (1997). Perceiving and Remembering: Category Stability, Variability and Development. In K. Lambert, & D. Shanks (Eds), *Knowledge, Concepts, and Categories* (pp. 161-195). MIT Press.
- Sottile, Z. (2022). California Bees Can Legally Be Fish and Have the Same Protections, a Court Has Ruled . *CNN*.
- Strevens, M. (2000). The Essentialist Aspect of Naive Theories. Cognition, 74, 149-175.
- Strevens, M. (2001). Only Causation Ratters: reply to Ahn et al. Cognition, 82, 71–76.
- Strevens, M. (2019). Thinking Off Your Feet. Belknap Press.
- Tabossi, P. (1982). Sentential Context and the Interpretation of Unambiguous Words. *Quarterly Journal of Experimental Psychology*, *34*, 79-90.
- Tabossi, P., & Johnson-Laird, P. (. (1980). Linguistic Context and the Priming of Semantic Information. *The Quarterly Journal of Experimental Psychology*, 32(4), 595-603.
- Taylor, P., & Vickers, H. (2017). Conceptual fragmentation and the rise of eliminativism. *European Journal for Philosophy of Science*, 7(1), 17-40.
- Tversky, A. (1977). Features of Similarity. *Psychological Review*, 84(4), 327-352.
- Van Overschelde, J., Rawson, K., & Dunlosky, J. (2004). Category norms: An updated and expanded version of the Battig and Montague (1969) norms. *Journal of Memory and Language*, 50, 289-335.
- Vicente, A. (2018). Polysemy and Word Meaning: An Account of Lexical Meaning for Different Kinds of Content Words. *Philosophical Studies*, 175(4), 947-968.
- Vicente, A., & Falkum, I. (2017). Polysemy. In Oxford Research Encyclopedia of Linguistics. Oxford University Press.
- Vicente, A., & Martínez Manrique, F. (2016). The Big Concepts Paper: A Defence of Hybridism. *British Journal for the Philosophy of Science*, 67(1), 59-88.
- Visser, A. (2023). Genericity Generalized. *Philosophical Studies*, 180, 703–723.
- Voorspoels, W., Storms, G., & Vanpaemel, W. (2013). Idealness and Similarity in Goal-Derived Categories: A Computational Examination. *Memory & Cognition, 41*, 312-327.

- Voorspoels, W., Vanpaemel, W., & Storms, G. (2011). A Formal Ideal-Based Account of Typicality. *Psychonomic Bulletin and Review*, 18, 1006-1014.
- Weiskopf, D. (2009). The plurality of concepts. Synthese, 169, 145-173.
- Weiskopf, D. (2020). Anthropic Concepts. Noûs, 54(2), 451-468.
- Whitney, P., McKay, T., Kellas, G., & Emerson, W. (1985). Semantic Activation of Noun Concepts in Context. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(1), 126-135.
- Willemson, P., & Reuter, K. (2021). Separating the Evaluative from the Descriptive: An Empirical Study of Thick Concepts. *Thought: A Journal of Philosophy*(10), 135-146.
- Williamson, T. (1994). Vagueness. London and New York: Routledge.
- Wilson, M. (2006). Wandering Significance. Oxford: Claredon Press.
- Wittgenstein, L. (2009). Philosophical Investigations, 4th Edition revised. Wiley-Blackwell.
- Yee, E., & Thompson-Schill, S. (2016). Putting Concepts into Context. *Psychonomic Bulletin & Review*, 1015-1027.
- Yeh, W., & Barsalou, L. (2006). The Situated Nature of Concepts. *The American Journal of Psychology*, 119(3), 349-384.
- Zadeh, L. (1965). Fuzzy Sets. Information and Control, 8(3), 338-353.