



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/209221/>

Version: Published Version

Article:

Margolis, E. and Laurence, S. (2023) Making sense of domain specificity. *Cognition*, 240. 105583. ISSN: 0010-0277

<https://doi.org/10.1016/j.cognition.2023.105583>


Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Making sense of domain specificity

Eric Margolis^a, Stephen Laurence^{b,*}

^a Dept. of Philosophy, 1866 Main Mall, E370, University of British Columbia, Vancouver, BC V6T 1Z1, Canada

^b Department of Philosophy and Hang Seng Centre for Cognitive Studies, University of Sheffield, 45 Victoria Street, Sheffield S3 7QB, United Kingdom

ARTICLE INFO

Keywords:

Domain specific
Domain general
Rationalism
Empiricism
Nativism
Innateness

ABSTRACT

The notion of domain specificity plays a central role in some of the most important debates in cognitive science. Yet, despite the widespread reliance on domain specificity in recent theorizing in cognitive science, this notion remains elusive. Critics have claimed that the notion of domain specificity can't bear the theoretical weight that has been put on it and that it should be abandoned. Even its most steadfast proponents have highlighted puzzles and tensions that arise once one tries to go beyond an initial intuitive sketch of what domain specificity involves. In this paper, we address these concerns head on by developing an account of what it means for a cognitive mechanism to be domain specific that overcomes the obstacles that have made domain specificity seem so problematic. We then apply this understanding of domain specificity to one of the key debates that it has figured prominently in—the rationalism-empiricism debate concerning the origins of cognitive traits—and introduce several related theoretical notions that work alongside domain specificity in helping to clarify what makes a view more (or less) rationalist. This example illustrates how the notion of domain specificity can, and should, continue to play a central role in ongoing debates in cognitive science.

1. Introduction[†]

The question of whether any cognitive mechanisms are domain specific is at the centre of some of the most important debates in cognitive science. Although the term “domain specificity” wasn't commonly used in the earliest days of cognitive science, the idea behind it goes back at least as far as Chomsky's discussions of the origins of knowledge of language and his critique of the then dominant accounts of the origins of language, which took language to be acquired solely on the basis of general-purpose learning mechanisms that are responsible for acquiring knowledge across many areas of cognition (Chomsky, 1965, 1967, 1972).¹ Chomsky's highly influential alternative was an approach in which language acquisition is grounded in a special-purpose faculty

for acquiring knowledge of language that incorporates principles distinctive to language. Chomsky also went on to speculate that cognitive development pertaining to other capacities might take a similar form, while noting that this general approach, which postulates distinct specialized cognitive mechanisms for acquiring different cognitive capacities, can be traced back to Descartes and Leibniz and other rationalist philosophers, just as competing approaches that suppose that cognitive capacities largely originate in the very same general-purpose mechanisms can be traced back to Hume and other empiricists. For Chomsky, one of the major advances offered by the scientific study of language in the generative tradition was that it illustrated how we can move beyond vague speculations about the origins of knowledge by formulating competing hypotheses with enough precision that empirical

* Corresponding author.

E-mail addresses: eric.margolis@ubc.ca (E. Margolis), S.Laurence@sheffield.ac.uk (S. Laurence).

[†] This article was fully collaborative; the order of the authors' names is arbitrary.

¹ To the best of our knowledge, the earliest discussion in cognitive science to explicitly employ the terms “domain specificity” and “domain generality” in a contrasting way was Keil (1981), which argued that we should think of cognitive development in terms of constraints on the acquisition of knowledge. Inspired, in part, by Chomsky's work in linguistics, Keil suggested in this important paper that “[t]here are not merely a few general constraints that apply without prejudice to all instances of knowledge acquisition... Rather, there are sets of highly restrictive constraints specifically and uniquely tailored to various cognitive domains...” (p. 198).

data and argumentation can be brought to bear on the traditional debate between rationalism and empiricism.²

The notion of domain specificity, in one form or another, soon came to play a central role in a number of related debates in cognitive science, including debates about whether aspects of perception and cognition are modular (e.g., Carruthers, 2006; Fodor, 1983; Sperber, 1994), about the existence and prevalence of psychological adaptations (e.g., Barrett, 2015; Tooby and Cosmides, 1992), about the basis of cognitive disorders and neuro-atypical cognition (e.g., Kamps et al., 2017; Leslie and Thaiss, 1992), about the basis of category-specific deficits (e.g., Mahon, 2022; Santos and Caramazza, 2002), about the existence and nature of distinctively human forms of intelligence (e.g., Mithen, 1996; Spelke, 2003), and about the origins of cross-cultural universals (e.g., Atran and Medin, 2008; Boyer, 2022), among others.

Yet, despite the widespread reliance on domain specificity in recent theorizing in cognitive science, this notion remains elusive. Some critics have claimed that the notion of domain specificity can't bear the theoretical weight that has been put on it and that it ought to be abandoned (e.g., Prinz, 2006; Woodward and Cowie, 2004). And even among theorists who make heavy use of the notion of domain specificity, there is a sense that it is hard to pin down once one tries to go beyond an initial intuitive idea. For example, in describing their influential edited volume of papers on domain specificity, Hirschfeld and Gelman remark that it was based on a conference that "was an attempt to discover if the notion of domain specificity could be discussed profitably (even intelligibly!) across disciplinary lines" (Hirschfeld and Gelman, 1994a, p. viii).

Efforts to provide a more detailed understanding of domain specificity have also notably given rise to a number of puzzles and difficulties that have proven difficult to resolve. For example, while many theorists take domain specificity to apply to both cognitive mechanisms and informational states (e.g., Chomsky, 1980, 1984; Hirschfeld and Gelman, 1994b), others take the application to informational states to be essentially trivial and uninteresting, as in this passage from Fodor:

... information is ipso facto specific to the domain that it is information about. The information that cows have horns is specific to cows. The information that everything that exists is spatially extended is specific to everything that exists; the information that cats scratch is specific to cats; and so on. This notion of the domain specificity of information is, patently, of no use to anybody. (Fodor, 2000, p. 58).

And, even in relatively simple cases, it can be unclear how to apply the distinction between domain specificity and domain generality. Notice, for instance, that an inference rule like *modus ponens* seems to be domain specific in some ways and domain general in others. Fodor calls attention to the unclarity about whether a rule like this counts as domain specific or domain general, remarking that *modus ponens* only applies to arguments with a single precise form (*if P then Q; P; therefore Q*), which makes it seem domain specific, but that at the same time it doesn't matter what the *Ps* and *Qs* refer to, which makes it seem domain general (Fodor, 2000, p. 59). The various options here—*modus ponens* is domain specific, *modus ponens* is domain general, and *modus ponens* is both domain specific and domain general—seem equally unattractive. These sorts of puzzles and difficulties—and the further difficulties that, as we will see, arise from tempting solutions to them—can make the notion of domain specificity seem inherently problematic.

We will argue, however, that there is no reason to abandon the notion of domain specificity. The puzzles and problems that have surrounded the notion of domain specificity are not intractable; they can be resolved in a consistent and theoretically satisfying way. Properly

² Following Chomsky, we will refer to the debate concerning the origins of psychological traits as the *rationalism-empiricism* debate. Others use the term *nativism-empiricism* debate instead; for these purposes *rationalism* and *nativism* are merely stylistic alternatives.

understood, domain specificity is a perfectly coherent and vital theoretical construct that can—and should—continue to play a central role in a variety of debates that are at the very heart of cognitive science.

2. What is a domain? What makes a cognitive mechanism specialized for a given domain?

To begin, we need to say something about what a domain is. A domain—or as we will often say, a *content domain*—is best thought of as a subject matter (Fodor, 1983). It is the subject matter that a psychological structure is directed at.

Being directed at is a relation much like *represents*, possessing the same idiosyncratic properties that are characteristic of a whole family of related notions, including *being about* and *having intentionality*. One of these properties is being perspectival. This means that a content domain shouldn't be understood merely as a collection of entities. Built into the very idea of a content domain is that there is a way that the entities in that content domain are to be construed. As a consequence, two psychological structures could be directed at different content domains even if the two content domains contained exactly the same entities. To use a well-worn example, even if every creature with a heart is also a creature with kidneys, a mechanism that is specialized for representing creatures with hearts would be directed at a different content domain than a mechanism that is specialized for representing creatures with kidneys. Or, to use another example, a mechanism could be specialized for representing three-angled closed polygons as opposed to three-sided closed polygons, even though all triangular polygons are trilateral polygons, and vice versa. The perspectival nature of *being directed at* also means that the same entity can belong to many different content domains. For example, the same entity can be in the domain of physical objects, the domain of animals, the domain of agents, and so on. This entity would be represented in different ways and for different purposes by different domain-specific cognitive mechanisms directed at these different content domains.³

A second and related feature of the relation *being directed at* is that subject matters needn't correspond to objective categories discovered by science and may even involve things that don't actually exist. Just as a story could have unicorns as its subject matter or have as its subject matter how things might have gone if the dinosaurs hadn't become extinct, a cognitive mechanism could have as its content domain mythical or fictional creatures or be concerned with counterfactual events. For much the same reason, the content domain which a cognitive mechanism is directed at could diverge from the categories that are recognized and investigated by science even when the domain involves real as opposed to fictional entities. For example, while standard biological taxonomies don't recognize categories like *tree* or *fish* (because these categories don't involve groupings of organisms that include all of the descendants of a common ancestor), a cognitive mechanism might well range over these categories and have them as part of its content domain. This would be true, for example, of a mechanism that is responsive to the biological realm as it is conceptualized in everyday

³ Some proponents of domain specificity also distinguish between a mechanism's *proper domain* and its *actual domain* (Sperber, 1994). This distinction depends on a further commitment in which domain-specific mechanisms are seen as typically being cognitive adaptations. Given this assumption, the proper domain would be the domain that the mechanism evolved to deal with, and its actual domain (or domains) would be the domain(s) that it currently deals with. Part of the point of making this distinction is that the current environment may differ significantly from the environment in which the domain-specific mechanism evolved, and consequently there may be current items that the mechanism is responsive to (which are part of the mechanism's actual domain) even though they wouldn't have been present when the mechanism originally evolved and so didn't figure in the selection pressures responsible for the mechanism's adaptive functional structure (and so are not part of the mechanism's proper domain).

thinking (folk biology) as opposed to how it is conceptualized in scientific biology (Medin and Atran, 1999).

One further point stemming from the perspectival nature of content domains is worth commenting on, as it involves a potential source of resistance to domain specificity. In the course of introducing the notion of a content domain in *The Modularity of Mind*, Fodor comments that he will “also assume that we can make some sense of individuating content domains independent of the individuation of cognitive faculties, since if we cannot the question whether the operation of such faculties cross content domains doesn't arise” (Fodor, 1983, p. 13). It might seem that this line of thought implies that content domains need to be natural groupings of entities, where the naturalness is a product of some kind of independent fact about the world. Many, however, would be suspicious of any claims that the world comes pre-packaged into content domains in this way, and so would be suspicious of any notions that depended on such assumptions.

In response, we want to say that no questionable metaphysical assumptions of this sort need to be made. Content domains aren't delineated by mind-independent facts about what constitutes a natural category. They are simply a reflection of human psychology, in particular, the cognitive mechanisms that are directed at them. This should already be apparent from the fact that there could be a content domain of mythical creatures—creatures which do not and never have actually existed—or a content domain for a biological category that doesn't correspond to one that is recognized in scientific taxonomies in biology. But even in cases where a content domain does include entities that are recognized by science, the point remains that content domains are not individuated by some objective standard of naturalness. They are individuated by facts about what kinds of categories *minds like ours treat as natural*. Some categories are clearly not categories that our kinds of mind take to be natural—arbitrary disjunctive categories, for example.⁴ This is why cows, shoelaces, exoplanets, and musical instruments do not form a single content domain, since minds like ours don't take these entities to form a natural category.⁵ In contrast, other categories are clearly ones that our kinds of minds do take to be natural, for instance, categories corresponding to concepts that are expressed by individual words in a speaker's language. This is why a category like *tools* constitutes a single content domain even though it consists of many different sub-kinds (hammers, screw drivers, saws, rulers, and so on) that may otherwise be quite different from one another.

Two further potential concerns about content domains deserve mention. The first is whether some subject matters are too narrow or

insignificant to count as content domains. For example, assuming that there is unified body of information and set of motor schema associated with a cognitive mechanism that is responsible for our being able to shake hands with others, these structures clearly pertain to handshaking. But should we really say that there is a handshaking domain? One might think that this would be extending the notion of a domain too far. The second concern is tied to the character of the developmental processes involved in acquiring a cognitive mechanism that deserves to be called domain specific. Should idiosyncratic cognitive structures that are not universal and that require extensive input to develop be considered domain specific in the same sense as specialized mechanisms that are human universals and have a highly regular development trajectory? For example, should chess expertise be thought of as domain specific in the same sense as the domain-specific language faculty posited in Chomsky's seminal work? What about other esoteric forms of expertise, like expertise in slight-of-hand magic or exhaustive knowledge of medieval heraldry?⁶

As we see it, these concerns highlight the fact that while domain specificity plays a key role in some of the most important debates in cognitive science, the property of domain specificity, on its own, does not guarantee that a psychological structure is interesting or significant. Mundane or idiosyncratic traits can sometimes be domain specific. It is easy to be misled on this point by the fact that the prototypical examples of domain-specific mechanisms that have been proposed and that have been at the centre of important debates in cognitive science have been complex and richly structured universal cognitive mechanisms that are of unquestionable theoretical significance and that have been seen by rationalists to have substantial innate components.

Consider the handshaking example. The set of resources that the mechanism in question makes use of has a very narrow focus and isn't particularly theoretically interesting. But it nonetheless does seem to be specific to a particular domain. Other examples of equally fine-grained domains suggest that having a narrow domain is perfectly compatible with domain specificity. In a related discussion, Samuels (2006) notes that edge detectors in low-level vision seem unproblematically domain specific even if they are not as exciting as larger and more encompassing domain-specific mechanisms, such as the full visual system or a mechanism for representing and reasoning about mental states. We agree. As we see it, the fact that a mechanism is narrow shouldn't matter to whether it can be domain specific. Nor should it matter if it is mundane or peculiar.⁷ Similar points apply to capacities that develop only in certain individuals who have the drive and experience to acquire a narrow form of expertise. Once it is recognized that a domain-specific mechanism can be domain specific without having all of the features associated with the language faculty and other prototypical examples of domain specificity, there is no reason to suppose that something like chess expertise—which clearly does seem to be directed at a single content domain, after all—shouldn't count as domain specific despite the fact that it involves idiosyncratic expertise acquired through a long

⁴ Paradigmatically unnatural categories would also include categories like *grue* and *bleen* (Goodman, 1954). By definition, an object is *grue* if it is either green and examined before time *t*, or else blue and examined after time *t*; and an object is *bleen* if it is either blue and examined before time *t*, or green and examined after *t* (Goodman, 1954). In a related discussion about what makes a concept natural, Landau and Gleitman (1985) note that Goodman's *grue* and *bleen* suggest that there may be a continuum of naturalness in which some concepts are more natural than others for most learners and that *grue* and *bleen* are near the end of the unnatural side of such a continuum.

⁵ Hypothetically, though, if we had different kinds of minds, then it too would count as a content domain. This would be the case, for example, if our minds contained a cognitive mechanism that processed information about all and only cows shoelaces, exoplanets, and musical instruments, treating them as a single kind of thing. Of course, our minds don't operate that way, making such a potential content domain seem bizarre and unnatural. While this particular content domain isn't one that is likely to be natural for any human population, there is bound to be a certain amount of variation in what content domains exist across cultures, populations, or even individuals. And undoubtedly some domains that are treated as natural for some thinkers will seem bizarre or unnatural to others who think differently. This variation makes sense in light of the fact that what counts as a natural category is not intrinsic to the category but rather is a reflection of what given individuals' psychological processing mechanisms treat as unified categories and are directed at.

⁶ We haven't yet said what exactly makes it the case that the language faculty is domain specific (assuming that it is); we will come to that soon. For the moment, we are just raising the question of whether the notion of domain specificity needs to be restricted to rule out mechanisms that are associated with narrow or mundane capacities or capacities that are contingent on particular types of interests and experiences that are not widely shared.

⁷ Briefly commenting on a potential mechanism with a very narrow and peculiar domain—a “mechanism for reasoning about camels in London on Tuesdays”—Samuels notes that the problem with such a mechanism isn't that it wouldn't be domain specific. It's that it is implausible to suppose that there would be a genuine cognitive mechanism with this domain (2006, p. 50).

history of exposure to chess.⁸

At this point, it will be useful to introduce a further piece of terminology. We need a way of collectively referring to the concepts that are associated with a given content domain through being directed at that domain. We will use the term *conceptual cluster* for this purpose. Suppose, for example, that there is a learning mechanism that is specialized for acquiring animal concepts. Given the right types of experience, it generates specific animal concepts—ZEBRA, BOA CONSTRICTOR, FALCON, GREAT WHITE SHARK, and so on. Taken together, all of these concepts would constitute a conceptual cluster which is directed at, or has as its subject matter, the content domain *animals*. Having this terminology in place helps to keep clear whether one is referring to the subject matter (content domain) or to the psychological structures that are directed at the subject matter (in this case, the conceptual cluster that is directed at the content domain).⁹ A similar issue regarding potential unclarity also arises for more complex informational states, where psychologists often use the term *body of knowledge*. This term could be taken to refer to the subject matter (what it is knowledge of) or to the psychological states that encode and process information pertaining to this subject matter. To be clear, when we speak of something as a body of knowledge, we are referring to psychological states that are directed at a subject matter, rather than the subject matter itself. For example, a body of knowledge that is specific to physical objects and core physical interactions between such objects is a psychological structure that is directed at the content domain *physical objects*. To sum up the terminology so far: A content domain is a subject matter. In contrast, conceptual clusters and bodies of knowledge aren't subject matters; they are psychological structures

which are directed at particular subject matters.¹⁰

Having clarified what a content domain is and having addressed some important concerns regarding domains and domain specificity, we are now ready to consider the question of what makes for domain specificity or domain generality directly. There are several related issues that need to be sorted out (which we will turn to in the next section), but to a first approximation, domain specificity is a matter of being directed at a particular domain, whereas domain generality is a matter of being directed at a number of distinct domains. When a cognitive mechanism is domain specific, what makes it the case that it is directed at a given content domain? One common answer to this question is that it is a matter of the input to the mechanism. Carruthers (2006) sees domain specificity in these terms, distinguishing a mechanism's input from other information it may access in the course of its operations. Input for Carruthers is understood in terms of what "turns on" the mechanism. For example, supposing a cognitive mechanism were only turned on by linguistic input, then this mechanism would be considered to be a domain-specific mechanism that is directed at the content domain of language.

The idea that input is what makes a domain-specific cognitive mechanism be directed at a given content domain isn't the only possibility, however. Other theorists emphasize the nature of the computations that take place within the mechanism (e.g., Cosmides and Tooby, 1994; Gallistel, 2003). Consider, for example, a cognitive mechanism that Cosmides and Tooby have proposed which is dedicated to determining whether those receiving benefits in social exchanges are entitled to them (often referred to as a *cheater detection module*). The proposed mechanism is taken to employ distinctive processes that are specially tailored to determining the legitimacy or illegitimacy of received benefits in social exchanges. Yet the input to this mechanism can be highly diverse. The relevant benefits might involve financial gain, admittance to a desirable school system, the right to drive someone's car, or any number of other things, and the legitimacy of such benefits might turn on a huge variety of factors. This way of determining what makes a cognitive mechanism specific to a particular domain allows for domain specificity in cases where the mechanism may have diverse inputs but is nonetheless directed at a particular content domain in virtue of the fact that the processing mechanism is specifically tuned to processing content from the content domain that it is directed at.

There is also a third factor that should be considered regarding what makes a domain-specific cognitive mechanism be directed at a given content domain, one that has been largely neglected in the literature on domain specificity. This has to do with the *output* of the mechanism. Let's consider again a hypothetical domain-specific mechanism that is solely devoted to acquiring concepts of animals. Arguably, a key feature that makes such a mechanism directed at the content domain of animals is the fact that the output of this mechanism is the conceptual cluster that is directed at the content domain of animals. One of the advantages of using output to determine what makes a cognitive mechanism domain specific is that in many cases we may not have much information about how a mechanism works or exactly what type of input it is restricted to, and it is more straightforward to simply consider its output. For example, if a learning mechanism just produces representations of faces, then it is specialized for the content domain *faces*, and we do not need to

⁸ This means that there is the potential for an unbounded number of domain-specific mechanisms corresponding to the many different areas of narrow expertise that might be acquired. A reviewer has suggested that this might be a good reason to distinguish between at least two types of domain specificity, one for domain-specific mechanisms that are innate and one for domain-specific mechanisms that implement learned expertise. We don't see the notion of domain specificity as applying differently in these two types of cases. It is just that in one case the domain-specific mechanisms are innate, and in the other they aren't. So, the mechanisms fall into two distinct categories not by involving different types of domain specificity but in light of combining domain specificity with a further feature (being innate or being learned). (See Section 5 for discussion of how domain specificity relates to the rationalism-empiricism debate, where a prominent role is given to the innate domain-specific mechanisms). At the same time, it should also be noted that it is an empirical question whether any particular narrow form of expertise is implemented by a domain-specific mechanism. In such cases, the expertise is specific to a domain, but many theorists hold that these forms of expertise can be implemented by domain-general mechanisms (see footnote 22 for an example of this type of view).

⁹ A reviewer has asked what makes a domain conceptual and whether a domain should be considered conceptual when tacit knowledge is involved. The distinction between domains and conceptual clusters is relevant here. Domains themselves are simply subject matters and so are neutral with respect to the type of psychological entities that might be directed at them. All variety of psychological structures—including conceptual clusters and cognitive mechanisms involving concepts, but also nonconceptual representations and cognitive mechanisms operating over nonconceptual representations—can be directed at a domain. For this reason, we prefer to avoid the term "conceptual domain". But, one might still wonder what it is that makes a cognitive mechanism (or a representation) conceptual. This is an important question, but it's one that we take to be completely independent of questions about what domains are. And the question of what makes a mechanism or representation conceptual is too controversial and too complex to address in this paper. For discussion of different ways of drawing the conceptual/nonconceptual distinction, see Laurence and Margolis (2012).

¹⁰ Nothing turns on adopting this terminology apart from being clear about whether one is referring to a subject matter or to psychological structures that are directed at a subject matter. However, we should note that other theorists sometimes use the same term for both and let the context indicate which meaning is intended. For example, Hirschfeld and Gelman (1994b) use the term "domain" to refer to a subject matter but also use the same term to refer to "a body of knowledge that identifies and interprets a class of phenomena assumed to share certain properties and to be of a distinct and general type" (p. 21). We prefer to avoid this dual use of the term "domain" in order to avoid potential confusions it may generate.

know whether its internal computations are uniquely suited to faces in order to see that the mechanism is domain specific.¹¹

In our view, rather than trying to decide which of these three factors—distinctive input, specialized internal processes, or distinctive output—is most important, or trying to distinguish different senses of domain specificity linked to these factors, it is better to understand domain specificity as a function of all three factors. In particular, while we think that any of the three factors suffices for domain specificity, we see domain specificity as involving all three factors.¹²

Take, for example, Chomsky's classic proposal that there is an innate language faculty, an innate domain-specific mechanism for acquiring natural language syntax. By hypothesis, this mechanism produces just one thing—a grammar that specifies the syntactic properties of the local natural language. If a young learner is exposed to more than one natural language, their language faculty may produce further grammars for each of these languages, but it can't do much else. It can't acquire knowledge of the rules of chess; it can't figure out how to navigate through a maze; it can't help you balance your chequebook. It can't even produce a grammar for some other type of system of communication—a “language”—whose structural properties substantially deviate from those of human natural languages. This is because it is directed at languages that conform to the principles of Universal Grammar. Clearly, then, this mechanism is quite limited regarding its output. It is also limited regarding its input. The language faculty, on this proposal, is selective regarding the information it is responsive to and uses when forming a grammar. It doesn't respond to sounds in general or even more narrowly to the vocalizations emitted from other individuals. Its input consists of *linguistic* expressions (words, phrases, sentences), which it represents specifically as linguistic data. Finally, the language faculty exploits this incoming information in a distinctive manner. On one such proposal, the language faculty embodies a set of parameters each of which has just a few options regarding some critical syntactic property. For example, a parameter may determine whether a language is *head-initial* or *head-final* (the head of a phrase being the word that establishes the phrase's syntactic category—for example, the verb in a verb phrase). In head-initial languages, the head appears before its complements; in head-final languages, it appears after its complements. The point is that a mechanism that incorporates a number of parameters of this kind, which are specific to structural features of natural language, is uniquely suited to acquiring languages that conform to Universal Grammar and hopelessly unsuited to doing anything else. So, this prototypical case of a domain-specific mechanism clearly involves all three factors. The language faculty is specialized for this one domain because of the type of input it relies on, the way it processes this input, and the type of output it can produce.

¹¹ Such a mechanism might employ a form of statistical analysis that could be equally used in mechanisms that are provided with different types of content. Nonetheless, as we see it, if it were part of an overall cognitive architecture in which it was positioned to only receive input involving facial stimuli and, as consequence of this arrangement, delivered output of just one type of content (facial representations), it would still count as a domain-specific mechanism.

¹² In other words, domain specificity covers various types of cases in a way that is analogous to a preference rule system. With preference rule systems, several criteria may be associated with a term, where none of these must absolutely be satisfied—any one of the criteria will do, although the satisfaction of all of the criteria makes for more stereotypical instances (Jackendoff, 1983). As an example, Jackendoff, notes that the term “climbing” covers the condition of moving upward and of moving with grasping motions involving some effort. If both of these conditions are met (*Bill climbed up the mountain*), the situation involves climbing. If only one is met (*Bill climbed down the mountain* / *The snake climbed up the tree*), then the situation still involves climbing, though in a less stereotypical way. But if neither is met (*The snake climbed down the tree*), there is no climbing. The point of the analogy is that we are suggesting that the conditions that make a domain-specific mechanism be directed at a given content domain are such that paradigmatically all three of these conditions are met, but any of the three is sufficient.

3. Does domain specificity come in degrees?

The terms *domain specific* and *domain general* are typically used categorically. A given cognitive mechanism is simply said to be domain specific or domain general (without qualification). Nonetheless, there is reason to suppose that domain specificity and domain generality might be graded phenomena and so come in degrees, and some authors have highlighted the fact that a cognitive structure might be neither fully domain specific nor fully domain general.¹³ In this section, we will argue that there is a perfectly reasonable way of understanding domain specificity and domain generality as graded phenomena. So, it is possible for two cognitive mechanisms to both be domain specific, with one being more domain specific than the other. At the same time, however, what seems to be the most natural way of understanding how domain specificity and domain generality could be graded turns out to be highly problematic, and so the gradedness of these phenomena is another important source of potential confusion that needs to be sorted out.

Consider, for example, the origins of our understanding of numerical quantity. Some accounts posit an innate mechanism that is solely dedicated to representing and processing numerical quantities. This type of mechanism is widely seen as domain specific. By contrast, other accounts of the origins of our understanding of numerical quantity hold that numerical quantities are initially represented in more generic terms using a mechanism that is not solely dedicated to numerical quantities but that represents and processes a number of different types of quantities (or to use the term more typically used in this work, different types of *magnitudes*) in an undifferentiated manner. One mechanism of this type might be responsive to numerical magnitudes, spatial magnitudes, and temporal magnitudes, treating all of these as instance of a single kind without discriminating between them. A mechanism of this sort is hard to conceptualize for an adult human, since we clearly differentiate between these types of magnitudes.¹⁴ To get a handle on what it means for the mechanism to treat these different kinds of magnitudes in an undifferentiated way, we can suppose that it conflates inputs across these domains, for example, summing across spatial and numerical inputs in determining what size of a space-time-number magnitude should be associated with a given input. By contrast, a mechanism that treated these magnitudes in a differentiated way wouldn't do this—it would be capable of processing each of the three types of magnitudes but would treat each as a distinct kind. A general magnitude mechanism that represents these magnitudes in an undifferentiated way is widely thought to be more domain general than a mechanism that is solely dedicated to representing and processing numerical quantities. In fact, some theorists take proposals broadly along these lines to vindicate domain-general learning. For example, Newcombe (2002) uses a similar example to illustrate how, on her view, domain-specific mechanisms emerge from “domain-general starting points” (p. 398).

To take another example, consider two hypothetical learning mechanisms that might be involved in acquiring representations of norms. The first mechanism is solely dedicated to acquiring moral norms. The second mechanism is involved in the acquisition of a broader class of norms including both moral norms and conventional norms (for

¹³ In a discussion of moral emotions, for example, Nichols (2005) notes that “although the effects of these emotions on cognition are probably not [fully] domain specific, neither are they perfectly domain general”, using the term *domain diverse* to label such structures (p. 368).

¹⁴ For a model along these lines, in which spatial, temporal, and numerical quantity are represented indiscriminately within a single magnitude system, see Walsh (2003). See also Mix, Levine, and Newcombe (2016) for discussion of other dimensions that might factor into a generalized magnitude system and both Mix et al. and Leibovich et al. (2017) for accounts of how previously undifferentiated magnitudes of different types might become differentiated in the course of development.

example, norms about talking during a movie in a cinema, or wearing shoes in the house), treating them in an undifferentiated way as simply norms.¹⁵ Again, it seems natural to suppose that the second mechanism is more domain general than the first.

Why take the more general magnitude mechanism and the more inclusive norm acquisition mechanism to be more domain general than the mechanism that is solely dedicated to representing and processing numerical quantities and the moral norm acquisition mechanism? A natural answer is that this is because the first two have broader or more general content domains than the second two. This view has strong appeal. But it can't be right. We can see this by asking in what sense the first two have broader content domains than the second two, such that having broader content makes them more domain general. Is it that one mechanism applies to more actual things? This can't be right, since a mechanism dedicated solely to acquiring concepts for animals (or a mechanism for processing such concepts) wouldn't be more domain general if more animals existed. And clearly it wouldn't help to consider the number of *possible* objects in a content domain, since virtually any given content domain (including that of animals) might include any number of entities in other possible circumstances. What about the number of distinct *kinds* of things in the content domain? This won't help either, for much the same reason. If there were more kinds of animals, then such a mechanism might end up producing more concepts. But they would all still be animal concepts. The mechanism's content domain would still be that of animals, and so the mechanism wouldn't be any more domain general. Likewise, a mechanism dedicated solely to processing quantities corresponding to natural numbers would have an infinite domain (since there are infinitely many natural numbers), but having an infinitely large domain wouldn't make this mechanism domain general—or even any more domain general—than a mechanism that is dedicated to processing content pertaining to a finite domain. After all, it would remain directed to a single homogenous domain, just like the mechanism for acquiring concepts for animals. Clearly something has gone wrong.

To see what this is, we need to remember that content domains are not merely collections of entities but are essentially perspectival. So, the domain specificity or domain generality of a given cognitive mechanism shouldn't be taken to be a matter of the sheer number of entities it encompasses. Rather, a mechanism's being domain specific or domain general is a matter of the *number of distinct content domains it is directed at*. To a first approximation, a domain-specific mechanism is one that is directed at a single domain, whereas a domain-general mechanism is one that is multiply directed in that it is directed at a number of distinct domains (we will elaborate on this characterization in a moment).¹⁶

We have just argued that domain specificity is not a matter of the size of the domain. This is also supported by considering an example like the following. Compare a cognitive mechanism that is solely specialized for acquiring concepts about physical objects and their interactions (e.g., concepts such as *PHYSICAL OBJECT*, *PHYSICAL SUPPORT*, and *CONTAINMENT*) with a cognitive mechanism that is specialized for acquiring just concepts for tools.¹⁷ Clearly the first of these encompasses more entities since every

tool is a physical object while a great many physical objects aren't tools. Nonetheless, both mechanisms should be considered domain specific—and to precisely the same degree—because they are each specialized for acquiring concepts pertaining to just one content domain—the first towards the content domain *physical objects*, the second towards the content domain *tools*.

In contrast, domain-general learning mechanisms are *not* specialized in this way. A domain-general concept learning mechanism might be equally suited to the acquisition of concepts of physical objects, tools, numbers, emotions, and much else. What makes a domain-general learning mechanism of this sort domain general is the fact that it is directed at more than just a few domains (especially when, as in the example just mentioned, these are diverse domains that are not closely related in content) (see Fig. 1). Importantly, a domain-general mechanism of this sort is directed at these various domains not by collapsing them into a broader domain but in a differentiated way; it is *multiply directed*—directed at each of the various domains it concerns separately—by being successively directed at each of these domains when it is processing information pertaining to that content domain.¹⁸ A domain-general concept learning mechanism would be capable of acquiring concepts in a variety of domains not in virtue of properties that all these concepts have in common as members of a single larger content domain, but rather in virtue of properties that they each have that make them members of their respective different domains. When such a domain-general learning mechanism acquires concepts in the tool domain, it is directed at the content domain *tools*. When it acquires numerical concepts, this very same mechanism is directed at the content domain *number*. And so on for other conceptual clusters and their content domains.

While it seems clear that the domain specificity or domain generality of a given cognitive mechanism is not simply a matter of the sheer number of entities involved or the size of a given content domain, it nonetheless seems as though a theorist who posits a mechanism devoted solely to processing numerical quantities (as opposed to one that processes undifferentiated continuous magnitudes) will end up with a view that is, in some sense, more domain specific. How can we account for this? An important clue comes from a fact that Newcombe (2002) highlights in relation to these contrasting views—namely, that the less restrictive mechanism may well give rise to multiple more restrictive mechanisms during the course of development. This suggests that the greater domain specificity associated with the more restrictive mechanism is not a matter of this mechanism itself being more domain specific than the less restrictive mechanism. Rather, it stems from the fact that when a more restrictive mechanism is posited, it is likely to be accompanied by *additional* domain-specific mechanisms corresponding to the other more restricted mechanisms that are thought to derive from the less restrictive mechanism.¹⁹ What we have isn't a view that postulates a mechanism with a *greater degree* of domain specificity; it's a view that postulates a *greater number* of domain-specific mechanisms. In all likelihood, a theorist who is committed to the existence of a domain-specific mechanism that is solely devoted to processing numerical magnitudes will also be committed to the existence of one or more additional

¹⁵ See Sripada and Stich (2006) for some ideas about how a mechanism like this might work.

¹⁶ A domain-general mechanism will be directed at multiple domains in virtue of how it relates to the same three factors that determine the domain that a domain-specific mechanism is directed at. In particular, a domain-general mechanism will be directed at multiple domains in virtue of it taking input from these multiple domains, producing outputs in these multiple domains, and having a processing mechanism that is not specialized for processing content from any particular content domain.

¹⁷ To be clear, what we are saying here is that the cognitive mechanism that is specialized for acquiring concepts about physical objects and their interactions represents them as such. It doesn't merely acquire concepts that happen to pick out physical objects, representing them in other terms. It represents them *as* physical objects.

¹⁸ In their initial state, domain-general processing mechanisms will not yet treat inputs and outputs from particular domains as belonging to distinct domains but can be seen to be domain general in virtue of taking inputs from a wide range of domains and having a processor that is not specialized for a particular domain. Relatedly, there is a derivative sense of domain generality associated with a type of processing, as opposed to a processing mechanism, where a type of processing counts as domain general to the extent that it is not specialized for a particular domain.

¹⁹ In section 5, we will return to the question of the relative specialization of cognitive mechanisms in the context of the rationalism-empiricism debate. We will argue that additional factors beyond domain specificity are important determinants of the degree to which a proposed learning mechanism is more (or less) rationalist.




Fig. 1. Being directed at a single domain versus being multiply directed at a number of domains. (a) Domain specificity is not a matter of the size of the domain that a domain-specific mechanism is directed at. A cognitive mechanism that is specialized for acquiring physical object concepts and a cognitive mechanism that is specialized for acquiring just tool concepts are equally domain specific in that each is directed at a single domain. The first is specialized for acquiring concepts pertaining to physical objects in general (e.g., the concepts PHYSICAL OBJECT, PHYSICAL SUPPORT, CONTAINMENT), the second for acquiring concepts pertaining to tools (e.g., the concepts HAMMER, SAW). The fact that there are more physical objects than tools does not make the first of these mechanisms less domain specific than the second. (b) By contrast, a domain-general mechanism for learning concepts is directed at more than just a few domains and is successively directed at each of these as such when it processes information pertaining to that content domain. Such a mechanism doesn't treat tools, animals, emotions, etc. as all being instances of a single undifferentiated kind. Rather it treats the items in these domains as categorically different from one another, allowing it to acquire concepts in each of these domains.

domain-specific mechanisms for temporal and spatial magnitudes as well. In contrast, a theorist who is committed to a generic magnitude mechanism that is responsive to numerical, temporal, and spatial magnitudes in an undifferentiated way is less likely to suppose that the mind contains additional mechanisms that are specialized for temporal or spatial magnitudes as well.²⁰

Does this mean that domain specificity and domain generality shouldn't be understood as graded phenomena, such that it is possible for one cognitive mechanism to be more domain general (and less domain specific) than another? Not necessarily. These phenomena may still be graded. To see how this can be, we need to consider a variant on the sort of general magnitude system that we have just been discussing.

We have been understanding this mechanism in such a way that it is responsive to numerical, spatial, and temporal magnitudes and that it responds to instances of these different types of magnitudes in an undifferentiated manner. Now consider a different type of mechanism, which is similarly responsive to numerical, spatial, and temporal magnitudes, but in this case, the mechanism *does* differentiate between numerical, spatial, and temporal magnitudes, representing numerical magnitudes as numerical magnitudes and not merely as continuous magnitudes (and likewise representing spatial magnitudes as spatial magnitudes, and temporal magnitudes as temporal magnitudes). Such a mechanism should be seen not as being directed at the content domain of continuous magnitude, but rather as being multiply directed at three

separate content domains—the content domain of spatial magnitudes, the content domain of temporal magnitudes, and the content domain of numerical magnitudes.²¹ This is because, as was emphasized earlier, we need to keep in mind the perspectival nature of content domains. In this case, that means that the mechanism doesn't treat the various magnitudes it ranges over as being all of a single kind (for example, it doesn't treat these different kinds of magnitudes as interchangeable, or simply sum them in considering what magnitude to associate with a given input).²² Rather, it is directed at each domain as such. Arguably such a mechanism should still be considered to be domain specific to some degree. It is not as domain specific as a mechanism dedicated solely to processing numerical magnitudes. But neither is it fully domain general. A typical domain-general mechanism would be responsive not just to content domains for different types of magnitudes but to a variety of other types of content domains as well, such as content domains for artifacts, for emotions, for colours, for foods, and so on. So, one way to go would be to say that domain specificity and domain generality should be understood as graded phenomena, and that a cognitive mechanism

²¹ Moreover, these three needn't be represented as all being instances of a single higher-level category (e.g., as types of magnitude). In some cases, they might be, in which case the mechanism would be directed at three content domains which are also represented as being related to one another. But in other cases, it would just be directed at these different content domains without representing that they are related.

²² McClelland et al. (2010) similarly highlight how a domain-general system can treat different domains as distinct, processing content in one domain differently than content from a different domain. They note that while in principle connectionist neural network mechanisms can be domain specific,

... connectionist work has focused on generic constraints that foster the discovery of structure, whatever that structure might be, across a range of domains and content types [(Hinton and Salakhutdinov (2006); Rogers and McClelland (2004)]. Yet, despite using only domain-general constraints, the connectionist model of semantic learning ... can acquire domain-specific patterns of responding ... [and can] exploit different types of similarity among the same set of items in different contexts (e.g. taxonomically-defined similarity for biological properties, but a one-dimensional similarity space for judgments about size ...). (McClelland et al., 2010, p. 353)

²⁰ Or at least a theorist who holds a view like this is unlikely to suppose that the mind *initially* contains such additional mechanisms—like Newcombe, such a theorist may hold that subsequent development leads to the differentiation of different types of magnitude and the emergence of several more specialized mechanisms. In this way, different theorists can in principle agree about the number and domain specificity of the cognitive mechanisms that are ultimately acquired but disagree about the number and/or domain specificity of the cognitive starting points of the developmental processes involved. (Though, of course, theorists may also disagree about the number and domain specificity of the cognitive mechanisms that are ultimately acquired.)

that is responsive to numerical, spatial, and temporal magnitudes in a differentiated manner is one that is domain specific but to a lesser degree than a mechanism that is responsive solely to numerical magnitudes.²³ This view takes one cognitive mechanism to be more domain general than another to the extent that it has a higher degree of multi-directedness: that is, it is directed at a larger number of different domains (especially when they are diverse in content), being successively directed at each of these domains as such when it is processing information pertaining to that content domain. Likewise, one cognitive mechanism will be more domain specific than another to the extent that it is directed at a smaller number of closely related content domains, where being directed at only one content domain counts as being maximally domain specific.

We are suggesting that this is a perfectly cogent way of understanding how domain specificity and domain generality can be viewed as graded phenomena. It is worth noting, however, that it is not necessary to see them as graded phenomena. While taking domain specificity and domain generality to be graded along the lines that we have just suggested seems to be a perfectly reasonable option, an alternative would be to avoid talk of gradedness altogether and instead hold that there is a threshold or cut-off point at which a mechanism is no longer considered to be domain specific. In that case, a mechanism would be held to be domain specific in virtue of its being directed at a single content domain or a sufficiently small number of sufficiently closely related content domains. And a mechanism would be held to be domain general in virtue of its having a sufficiently high degree of multi-directedness, being differentially directed at a sufficiently large number of sufficiently diverse content domains. Typical domain-general mechanisms are seen as being directed at *many* different domains, but even being directed at more than just a few (especially when these are diverse content domains) is normally, and in our view reasonably, taken to suffice for being domain general.

4. Resolving Fodor's puzzles

We are now ready to return to the issues that we mentioned at the start of this paper in connection with Fodor's views on domain specificity. Let's begin with the question of whether it makes sense to take informational states to be domain specific or domain general. In discussing this question, Fodor remarked that while it might be coherent to take informational states to be domain specific, such claims are of no interest at all because they are simply trivial. As he put it, "information is ipso facto specific to the domain that it is information about"—that is, any informational state is invariably specific to whatever content the state expresses and hence there is nothing substantive to the claim that it is domain specific.

As we see it, Fodor is simply mistaken here. There is a perfectly interesting nontrivial sense in which informational states can be domain specific. It just isn't the sense that Fodor has latched on to. To see why Fodor's way of understanding the domain specificity of informational states is problematic, notice that there is really no way of understanding how domain specificity in Fodor's sense would contrast with domain generality when applied to informational states. What are the contrasting cases of domain generality supposed to look like? Are

informational states supposed to be domain general when they *don't* apply to their domains?²⁴

How then should the notions of domain specificity and domain generality be extended to informational states? In the previous section, we argued that a cognitive mechanism should be understood to be domain specific in virtue of its being directed at a specific content domain or a small number of related content domains. We propose that the same basic idea can be extended to cover informational states. An informational state, such as a body of knowledge, is also domain specific in virtue of being directed at a specific content domain or a small number of related content domains.

The value of this way of thinking about domain specificity is easiest to see with regard to relatively substantial bodies of knowledge that might be associated with cognitive mechanisms. Consider two hypothetical mechanisms, one of which is associated with a body of knowledge that is just about dangerous animals and the other of which is associated with a body of knowledge that is in part about animals in general, but also in part about artifacts, in part about numerical quantities, in part about geometrical categories, in part about norms, and so on. The first body of knowledge counts as domain specific in virtue of the fact that it is directed at the content domain *dangerous animals*, while the second counts as domain general in virtue of the fact that it is directed at a large and diverse range of different content domains. As with cognitive mechanisms, informational states are more domain general to the extent that they are directed at a greater number and more diverse range of content domains. The applicability of this type of account is obscured by examples like Fodor's which focus on highly circumscribed informational states, but once one considers larger and potentially more complex informational states, the parallel with cognitive mechanisms is no longer obscured.

Let's turn now to the second issue, concerning Fodor's discussion of the inference rule *modus ponens*. Recall Fodor's observation that such an inference rule seems to be both domain specific and domain general. It seems specific in that it "applies only to arguments with premises of the form" *if P then Q; P*. On the other hand, it seems general in that "it abstracts entirely from the content of the premises it applies to" and so it applies equally across "quite different domains (physics and literary theory, as it might be)" (Fodor, 2000, p. 60).

Fodor attempts to resolve this tension by suggesting that a mechanism that instantiates *modus ponens* would be neither domain specific nor domain general. He claims, instead, that domain specificity and domain generality only apply "to the way that information and processes interact" (Fodor, 2000, p. 60). As an illustration of this claim, he gives the example of a mechanism that uses an inference rule that instantiates *modus ponens* but only works for *modus ponens* inferences involving properties associated with the number 2. Given the premises *2 is F* and *If 2 is F, then 2 is G*, it draws the inference *2 is G*. The idea seems to be that such a mechanism would be specific to the content domain of the number 2. However, this doesn't actually address the puzzle Fodor started with, since a mechanism that *isn't* restricted to *modus ponens* inferences involving 2 (i.e., one that draws such inferences about *any Ps* and *Qs*) would still seem specific in one sense and general in another. Fodor's response does nothing to resolve this tension.

Our suggestion for how to resolve the puzzle about whether *modus*

²³ Crucially, it is also less domain specific than a mechanism that represents numerical, temporal, and spatial magnitudes in an undifferentiated way. While such a mechanism would apply to precisely the same magnitudes as the mechanism that differentiates between them, the undifferentiated mechanism treats all of these magnitudes in an undifferentiated way as being part of a single domain, while the differentiated mechanism treats them as belonging to three distinct domains.

²⁴ Note also that one might mimic Fodor's argument and hold that claims about cognitive *mechanisms* being domain specific are always trivial because "mechanisms are ipso facto specific to the domains that they process information about". Since this clearly isn't enough to show that any claim that a mechanism is domain specific is trivial, we should be wary of Fodor's parallel argument that domain specificity is always trivial as applied to bodies of information.

ponens is domain specific or domain general involves distinguishing what we call *functional specificity* (and *functional generality*) from domain specificity (and domain generality).²⁵ The distinction between functional specificity and functional generality is easy to conflate with the distinction between domain specificity and domain generality. But it is an important separate distinction which has to do with the range of functions a cognitive mechanism has, that is, the range of cognitive operations or computations it can perform, such as computing the similarity to a prototype, drawing inductive inferences, or rehearsing information in working memory. When a mechanism only has one kind of function or a small range of closely related functions, then we will say it is *functionally specific*, and when it has more than a small range of functions, especially when they are diverse functions, we will say is *functionally general*.²⁶ Crucially, the question of what range of functions a mechanism has is distinct from, and independent of, the question of what range of content domains a mechanism is directed at. This means that a functionally-specific mechanism needn't also be domain specific. A mechanism can be functionally specific and *domain general*. This is precisely what is involved in the case of a processing mechanism that instantiates modus ponens. Such a mechanism is functionally specific in that it can only perform one kind of cognitive operation (inferring Qs from premises of the form *if P, then Q* and *P*), but it is domain general in that it can perform this cognitive operation on content drawn from any content domain.²⁷ The distinction between functional specificity and domain specificity allows us to clearly see how modus ponens can seem to be both domain general and at the same time somehow specific, and so to see how Fodor's puzzle about modus ponens can be resolved.

Some might argue that functional specificity is a trivial notion—that any cognitive mechanism is functionally specific in that it is specific to whatever functions it performs. This concern is analogous to the charge made by Fodor that the notion of domain specificity is trivial when applied to informational states because any such state is specific to whatever content it happens to have. Our response to this charge of

triviality is analogous to what we said earlier. As before, no room is made for the needed contrast—in this case, an understanding of functional specificity that allows for the possibility of a functionally-general mechanism. This just goes to show that functional specificity isn't a matter of being specific to whatever function(s) a mechanism has. It is a matter of the number and variety of functions it has. A mechanism that instantiates just modus ponens will be functionally specific as it can only perform one type of cognitive operation.²⁸ Likewise for a dedicated face recognition mechanism, which is both functionally specific and domain specific.²⁹ By contrast, many types of cognitive mechanisms have been proposed which are meant to be functionally general, performing a highly diverse range of cognitive operations across a broad range of domains. Examples include general-purpose Bayesian learning mechanisms and neural networks that aim to provide the basis for many cognitive functions, such as general-purpose deep learning networks.³⁰

We have seen that a cognitive mechanism can be *functionality specific and domain general* (e.g., modus ponens, passive co-occurrence models of semantic memory), *functionally specific and domain specific* (a dedicated face recognition mechanism, a mechanism for learning just concepts of animals), and *functionally general and domain general* (a general-purpose artificial neural network). This just leaves the combination of functional generality with domain specificity (see Fig. 2). Could there be a mechanism that is *functionally general and domain specific*? This would be a mechanism that could perform a wide variety of different types of functions but only for one type of content—for example, only regarding things having to do with the domain of food. Such a mechanism might learn new food concepts (but no other kinds of concepts), categorize encountered stimuli as food or not (but not do any other kind of categorization), store facts about events involving food for recall (but not facts involving any other kind of information), draw deductive inferences concerning information about food (but not be capable of reasoning about anything else), and so on. Mechanisms of this type

²⁵ This distinction is similar to ideas raised in Barrett (2009), which examines the notion of domain specificity in connection with a commitment to an adaptationist perspective, and in Sperber (1994) and Carruthers (2006), which are primarily concerned with offering an account of what modules are. For our purposes, we can remain neutral as to whether any of the traits in question are adaptations or what exactly a makes a cognitive mechanism a module.

²⁶ Functional specificity, like domain specificity, may be understood either as graded notion or as one involving a threshold. If treated the first way, this would mean that cognitive mechanisms can be more or less functionally specific. If treated the second way, there would be a dividing line between mechanisms that are functionally general and those that are sufficiently restricted to count as being functionally specific. As before, nothing really turns on which of these two ways of talking is adopted.

²⁷ It may help in seeing the domain generality and functional specificity of modus ponens to consider a hypothetical general reasoning mechanism which subserves a very broad class of different kinds of propositional inferences. Such a mechanism would implement deductive, inductive, and other types of reasoning about propositional contents pertaining to any content domain. A mechanism of this type would be a paradigm of a domain-general cognitive mechanism, and an account that took the mind to contain only this one cognitive mechanism for doing all its reasoning at the propositional level would clearly be a domain-general theory. At the same time, this mechanism would be capable of performing a number of different functions corresponding to the different types of inferences it was able to draw. Accordingly, it would be both domain general *and* functionally general. Now suppose that we gradually strip away inference types that the mechanism was capable of mediating while leaving unchanged the range of contents that the mechanism was able to interact with as the full range of propositional contents. It seems that we could do this until we reach the point where only one type of inference was left—just modus ponens. At this point, the mechanism would be highly functionally specific, subserving just a single function (performing modus ponens inferences), but would remain just as domain general as the domain-general reasoning system we began with.

	domain specific	domain general
functionally specific	dedicated face recognition mechanism	modus ponens inference mechanism
functionally general	[possible but theoretically unattractive combination]	general-purpose artificial neural network

Fig. 2. The distinction between domain specificity and domain generality is distinct from and interacts with the distinction between functional specificity and functional generality, though one of the four possible combinations is theoretically unattractive (see the text for discussion).

²⁸ Other examples of functionally-specific domain-general mechanisms would include the class of semantic memory models that have been described as passive co-occurrence models. The mechanisms in these models employ functionally specific processes which essentially track co-occurrences in the input irrespective of the content involved (Yee, Jones, and McRae, 2018). Functionally-specific domain-general mechanisms have also been proposed for a variety of other facets of memory, as, for example, in Atkinson and Shiffrin's classic account of working memory (Atkinson and Shiffrin, 1968).

²⁹ Note that functional specificity is distinct from the second factor we mentioned in our characterization of what makes a mechanism domain specific. That factor involved the processes internal to the mechanism being distinctively suited to a given content domain. In functional specificity, the specificity of the mechanism is for a particular type of function, irrespective of whether this function is specialized for a particular content domain.

³⁰ Functionally-general domain-general mechanisms have played an important role in accounts in cognitive science from its earliest days (Newell, Shaw, and Simon, 1958; Newell and Simon, 1972) to the present (e.g., LeCun, Bengio, and Hinton, 2015).

would seem to be possible in principle, but it's hard to see why any organism would have such a mechanism. The problem with functionally-general domain-specific mechanisms isn't that they are impossible, it's just that they are theoretically unattractive. If a mechanism is capable of performing a very broad range of different functions, then it is unlikely that it would be restricted to performing those functions in a single domain. And if a mechanism is specialized for a particular domain, then it is also likely to be relatively specialized in terms of the set of functions that it can perform in that domain.

5. Domain specificity, domain generality, and the rationalism-empiricism debate

In the previous sections, we have sketched our account of domain specificity, shown how this account can resolve a range of puzzles and confusions surrounding the notion of domain specificity, and distinguished domain specificity (and generality) from functional specificity (and generality). In this final section, we look at domain specificity in the context of one of the key debates that it has figured prominently in—the rationalism-empiricism debate concerning the origins of psychological traits.³¹ As we will see, some of the tempting theoretical choices that generated problems which threatened to undermine the notion of domain specificity reappear in the context of this debate. In showing how domain specificity factors into the rationalism-empiricism debate, we will argue that it is critical to not ask domain specificity to do more work than it should. Domain specificity isn't the only resource that can be drawn upon to elucidate what is at stake in this debate; it should be seen as working together with several related notions to clarify what makes a view more (or less) rationalist.³²

Many theorists and commentators write as if the rationalism-empiricism debate were one and the same as the *nature-nurture debate*, taking rationalists to place more weight on nature (understood as the contribution of genes) and empiricists to place more weight on nurture (understood as the contribution of experience). Having linked these two debates in this way, it's not uncommon for such theorists to then go on to reject the rationalism-empiricism debate on the grounds that thinking about genes and the environment in these ways is ill-founded (see, e.g., Lewkowicz, 2011; Lerner, 2015). This is a mistake. Although there is little sense to the idea that one or the other of genes and the environment is more important regarding *any* trait, much less any psychological trait—a point that is widely recognized—this doesn't mean we should give up on the rationalism-empiricism debate. This is because the two debates shouldn't be identified in the first place; the rationalism-empiricism debate about isn't about genes versus the environment (Laurence and Margolis, 2016).

It is also sometimes said that rationalists are committed to the existence of innate traits (and with this, to the viability of the notion of innateness) whereas empiricists can avoid this problematic commitment through their emphasis on psychological traits being learned. But this is also a mistake. Rationalists and empiricists are equally committed to the existence of at least some innate psychological traits (those that form the ultimate basis of learning) and agree that many, if not most, psychological traits are learned (Gallistel, 2000; Keil, 1999).

What then is the disagreement about? Essentially, it comes down to a disagreement about what types of psychological traits are held to be innate and to form the innate basis on which other psychological traits

are learned (Margolis and Laurence, 2013).³³ To a first approximation, the distinctive feature of empiricist approaches is to hold that the stock of innate psychological structures that are fundamental to human learning is austere in that it largely consists in a body of sensorimotor representations and a relatively small number of domain-general learning mechanisms. Rationalist approaches generally agree that the stock of innate psychological structures contains these kinds of things but hold that it also contains quite a bit more, including innate concepts or other types of abstract representations and, more importantly, innate domain-specific cognitive structures that form the basis for a number of domain-specific learning mechanisms. We will refer to these additional types of innate structures, which are more characteristic of rationalist approaches, as *characteristically rationalist psychological structures*.³⁴ Rationalism-empiricism debates can play out regarding views about the mind in general (about the innate psychological structures that are involved in all types of learning and development), about a class of related psychological traits (such as different types of emotions, or different types of concepts), or about one particular psychological trait, as when Chomsky argued for a rationalist view specifically about natural language syntax. In general, though, we can say that a view is more rationalist (or less empiricist) when it postulates a greater number and greater variety of characteristically rationalist psychological structures to explain the origins of the cognitive traits it is concerned with. For present purposes—given our focus on domain specificity—what matters is that innate domain-specific psychological structures are a paradigmatic type of characteristically rationalist psychological structure. Other things being equal, a view will be more rationalist when it postulates a greater number and variety of innate domain-specific structures and mechanisms.³⁵

In most cases, learning will draw upon a variety of psychological resources, building off of previously learned cognitive mechanisms, representations, and bodies of knowledge. So when we are assessing a proposal regarding the way that a trait is learned, we need to think about not just the proximate mechanisms that account for its acquisition but the psychological structures that account for the origin of these proximate mechanisms, and likewise the psychological structures that are involved in the acquisition of these more fundamental psychological structures, and so on, going all the way back to whatever innate psychological structures are ultimately involved. To highlight the relevant innate psychological structures that are fundamental to an account of a learned trait, we will say that the learned trait *traces back to* these innate structures. To use the language example once more, on the sort of classic account associated with Chomsky where the language faculty embodies innate principles of Universal Grammar, knowledge of one's natural language may depend on many previously learned things but crucially traces back to innate domain-specific components in the language faculty.

Now given this account of the rationalism-empiricism debate, it may seem that a puzzle arises for our account of domain specificity. Consider

³³ For those who are apprehensive about the notion of innateness, this account can be formulated in a way that doesn't depend on this notion. The debate can instead be understood directly as being a debate about the collection of unlearned psychological structures that forms the basis for learning all other psychological traits (see Margolis and Laurence, 2013).

³⁴ Note that this doesn't mean that empiricists can never accept any such structures, only that these types of innate psychological structures are particularly characteristic of rationalist approaches to cognitive development.

³⁵ Since empiricists favour a more austere collection of innate psychological structures as the basis for subsequent learning, it will also be true that the greater the number of functionally-specific domain-general innate mechanisms an account is committed to, the less empiricist the account will be. However, since our aim in this section is to focus on the important role of domain specificity in this debate, rather than to provide a comprehensive account of all of the dimensions that play into characterizing the rationalism-empiricism debate, we can ignore this further point here.

³¹ As noted earlier, we are following Chomsky in using the terms *rationalism* and *rationalism-empiricism* debate for what are sometimes referred to as *nativism* and *nativism-empiricism* debate.

³² For a more detailed discussion of the rationalism-empiricism debate, see Laurence & Margolis (forthcoming).

a case like the one that we mentioned in Section 3, involving two accounts of how representations of particular moral norms are acquired. On one of these, the acquisition of such representations involves a learning mechanism that traces back to an innate domain-specific psychological structure that is specialized for representing moral norms as such. On the other, the acquisition of these representations involves a learning mechanism that traces back to an innate domain-specific psychological structure that is specialized for representing norms in general in an undifferentiated way (i.e., it represents various types of norms without distinguishing between them—moral norms, norms of etiquette, conventional norms, and so on). The difficulty for our account of domain specificity is that, on our account, these two learning mechanisms are *equally domain specific* even though the first seems more rationalist.³⁶

It is tempting to respond to this worry by claiming that the domain-specific learning mechanisms that these theories appeal to *aren't* actually equally domain-specific. Perhaps we should just say that the first (which, after all, traces back to an innate psychological structure with a narrow domain—that of *moral* norms) is more domain-specific than the second (which traces back to an innate psychological structure with a broader domain—that of *norms in general*). However, this way of addressing the puzzle is highly problematic. As we argued in Section 3, although it may at first look like we can focus on the comparative sizes of the content domains associated with different mechanisms to determine which is more domain specific, that approach just won't work. If there actually were more animals, or more kinds of animals, or there were more possible kinds of animals, none of this would make a mechanism for acquiring concepts of animals any less domain-specific. Even though these changes would mean that such a mechanism was directed at a domain that was in some sense broader, the mechanism would remain equally specialized in that it would still be confined to acquiring just concepts of animals. And a mechanism with an infinite domain (e.g., the natural numbers) not only isn't domain general in virtue of having an infinite domain (much less maximally domain general), but can be just as domain specific as one with a finite domain. The sheer size of a domain is just not relevant to how domain specific a mechanism that is directed at it is.

As we see it, the resolution of this puzzle doesn't turn on treating the mechanisms as involving different degrees of domain specificity—so our account of domain specificity can be retained as is. Rather, it turns on the relation between the *target content domain* (which a learning mechanism as a whole is directed at) and the content domain associated with the innate resource that the learning mechanism traces back to. In particular, it turns on how closely related or, as we will describe it, how closely *aligned* these two content domains are. We will argue that the puzzle can be resolved by taking alignment to come in degrees and holding that, in addition to the other factors that make an account more rationalist, *the greater the alignment, the more rationalist the account*. To see how this works, we will need to work through a few examples. For these purposes, we will return to the examples we first mentioned when asking whether domain specificity comes in degrees—mechanisms for acquiring moral norms and mechanisms for acquiring natural number concepts.

With the first of these, the domain that is targeted for learning is the domain of moral norms. On one type of account of how these are

acquired, the learning mechanism traces back to a psychological structure for representing moral norms as such and hence one that is directed at the same content domain, that of moral norms. In that case, the target domain that the learning mechanism is directed at and the content domain that is associated with the innate domain-specific resources that the learning mechanism traces back to are clearly closely aligned with one another—they are the very same domain. On a different type of account, however, the learning mechanism traces back to a psychological structure for representing not moral norms but norms in general and hence one that is directed at the domain of norms as such. On this account, the two domains—the target domain (which remains the same as on the previous account) and the domain associated with the innate resource that the learning mechanism traces back to—are still reasonably closely aligned, the first being a proper subset of the second. But the alignment is clearly less close than in the first account, where the two domains are identical. What makes the first account of the origins of moral norms a more rationalist account, then, is not that the innate resources it traces back to are more domain-specific, but rather that the domain that they are directed at is more closely aligned with the target domain of the learning mechanism than in the second sort of account. This resolves the puzzle of how one account can be more rationalist than another despite the fact that the two accounts trace back to equally domain-specific innate resources. The resolution is that there is a further dimension that makes one account more rationalist than another, namely, how closely aligned the content domain of a learned trait is with the content domain(s) of the innate domain-specific resources that the learning mechanism traces back to.³⁷

Looking at another example will help to make this clearer. Let's consider again how concepts corresponding to the natural numbers are learned, where children come to be able to represent the numerical quantities *one, two, three, four, five, six* (and so on). It will help for us to consider a wider range of possible accounts here than we did in the related discussion in Section 3.³⁸ One approach to this question is organized around a domain-specific learning mechanism based in part on an innate domain-specific system for representing approximate numerical quantity—the *approximate number system*. This mechanism is widely thought to be shared with other animals and to represent numerical quantity as such, not just nonnumerical properties that happen to be correlated with number (Odic and Starr, 2018). A signature feature of this system is that while it is capable of representing an open-ended number of different numerical quantities, its numerical representations represent in an approximate way, and so lack the precision of representations of particular natural numbers. As a consequence, the approximate number system on its own doesn't have the wherewithal to represent numerical quantities corresponding to the natural numbers. Still, researchers have thought that it can play a major role in a larger arrangement of cognitive mechanisms that underlie the acquisition of concepts for natural numbers (see, e.g., Dehaene, 1997; Spelke and Tsivkin, 2001). There are a number of ways to develop this general approach. Dehaene's proposal is that the limits of the approximate number system are transcended through a learning process which involves the use of language and other symbolic capacities in environments in which children are exposed to numerical words and to counting procedures. Putting aside the details, the thing that we want to highlight is the degree of alignment between the content domain of the

³⁶ In section 3, we suggested that one reason why an account of the first sort might seem more domain specific is that theorists who posit domain-specific learning mechanisms with finer-grained domains will often posit further domain-specific learning mechanisms in other related domains. This is true, but it does not fully resolve the puzzle here, since it will not always be the case, and so can't be the full explanation for why one account is more rationalist. One can readily imagine, for example, that a theorist might postulate an innate domain-specific mechanism for just moral norms and no further innate domain-specific mechanisms for non-moral norms (taking all other norms to be acquired via this mechanism together with domain-general learning mechanisms).

³⁷ Degree of alignment is not the only further factor that could be appealed to. We focus on it here since our interest is in the role of domain specificity in the rationalism-empiricism debate, and alignment is closely linked to and interacts with domain-specificity.

³⁸ Because our interest in the accounts that we will sketch is solely to illustrate and clarify the notion of alignment, we will leave out many of the details regarding how the models in these accounts are supposed to work and primarily highlight the most important domain-specific components involved in the models.

approximate number system, which the learning mechanism draws on, and the content domain of natural numbers, which the learning mechanism as a whole is directed at. The alignment, while not perfect, is quite substantial, accounting for the fact that such an approach provides a quite strong rationalist account of the origins of natural number concepts. The approximate number system is an innate domain-specific system that is directed at a domain involving numerical quantities even if it is not the same domain of numerical quantities that constitutes the target domain for the learning mechanism. For Dehaene, what makes learning concepts like SEVEN and TWELVE a real possibility is the fact that the approximate number system gives children an early intuitive sense that the world contains numerical properties and not just groups or collections that are larger or smaller in terms of their non-numerical properties (such as their area, spatial density, luminance, and so on).

Now consider a second approach. On this account, the domain-specific learning mechanism for acquiring natural number concepts is one that builds on an innate domain-specific mechanism for representing a wide variety of different types of magnitudes in an undifferentiated manner. According to this alternative, concepts corresponding to the natural numbers trace back to an innate domain-specific system for representing magnitudes in general, not just numerical magnitudes. Leibovich, Katzin, Harel, and Henik (2017) offer an account along these lines. In contrast with Dehaene's approach, which holds that children initially have a sense of numerical magnitude, Leibovich et al.'s proposal is that children initially have a sense of magnitude-in-general. On Leibovich et al.'s account, this is later augmented as children become more capable of individuating items and subsequently undergo a shift in what they attend to as a consequence of trying to interpret number words. Additionally, as children develop a greater capacity for cognitive control and inhibition, they come to differentiate correlations in their experience by detecting when these correlations no longer hold (especially when features like numerical quantity and area no longer correlate), setting them up to tease numerical magnitudes apart from other types of magnitudes. This learned capacity for representing numerical magnitudes can then scaffold the acquisition of concepts of natural numbers. Again, it's not the details that matter so much as the overall approach. In this case, the account includes a sense of number much like Dehaene's, but here this sense of number is itself learned and is part of a lengthy process that stands between the acquisition of natural number concepts (concepts like SEVEN) and the most critical innate domain-specific psychological structures that they trace back to—in this case, one for representing and processing elements in the content domain of magnitudes-in-general. Note that the innate domain-specific resource that this account traces back to is less well-aligned with the target domain of natural numbers. This explains why this account is less rationalist than one like Dehaene's, which is more closely aligned with this target domain.³⁹

A third type of account, briefly mentioned in Section 3, holds that the mechanism for learning concepts of natural numbers traces back not to an innate domain-specific system for representing magnitude-in-general nor to one for representing approximate numerical quantities in particular, but to a restricted class of magnitudes—spatial, temporal,

and approximate numerical magnitudes—representing the content domain comprised of these three types of magnitudes in an undifferentiated way. Such a system would be more aligned with the target domain of natural numbers than an account like Leibovich et al.'s yet less well-aligned with this target domain than an account like Dehaene's. These three accounts nicely illustrate both how alignment is a graded notion and how the degree of alignment between the target domain of the learning mechanism and the domain of the innate resources it traces back to is an important factor in determining how rationalist the account is (see Fig. 3a - 3c).

Finally, there is one further type of approach that it is instructive to consider. This is an account of the origins of natural number concepts that is *more* rationalist than an account involving the approximate number system, not less. The learning mechanism in this type of account is one that incorporates a small number system, an innate domain-specific system for representing the first few natural numbers: *one*, *two*, and *three* (see, e.g., Margolis and Laurence, 2008; Margolis, 2020). Like the other accounts we've mentioned, the learning mechanism would also involve further resources, including an ability to learn count words in language and an ability to put these words into one-to-one correspondences with collections. The small number system involved in this account gives learners a firm foothold for representing natural number magnitudes that is expanded as they learn to represent the numerical consequences of changes to small collections and learn to relate small natural number magnitudes to external symbols within the context of a learned counting procedure.⁴⁰ The innate domain-specific small number system is not more domain-specific than the approximate number system or the general magnitude system. What makes the account involving the small number system more rationalist than either of these other accounts is that the small number system is more closely aligned with the target domain of natural numbers (much as the approximate number system is more closely aligned with this target domain than the general magnitude system) (Fig. 3d).

In sum, much of the difference between rationalist and empiricist theories turns on the different weights they give to domain-specific cognitive structures in their accounts of the origins of psychological traits. But domain specificity should not be asked to do *all* the work in differentiating between competing accounts in the rationalism-empiricism debate. Instead, domain specificity should be seen as being complemented by related notions, like that of the degree of alignment between domains, which explains how one learning mechanism can be

³⁹ The contrast between these accounts also illustrates a different important feature of alignment. Alignment is not just a *graded* notion, but also a *relative* one. The same innate domain-specific resources can be more (or less) aligned with *different target domains* depending on the content domain that the learning mechanism which they are involved in is directed at. For example, if what is being learned isn't concepts like SEVEN but the concept of magnitude (in general), then the innate domain-specific resource that Leibovich et al.'s account appeals to would be *more aligned* with the target domain than the innate domain-specific resource that Dehaene's account appeals to, and this greater alignment would result in a more rationalist account of the origins of this concept, where the learning mechanism is more specialized for this domain than a learning mechanism whose only innate domain-specific resource was the approximate number system. The key point is that *being more specialized for* is a relative notion which depends on the target domain of the learning mechanism.

⁴⁰ The innate domain-specific mechanism in this account represents small numerical quantities as such and shouldn't be confused with other accounts that make use of domain-specific mechanisms that respond to small numbers of objects but not by representing their numerical quantity (Laurence & Margolis, 2005). For example, Spelke and Tsivkin (2001) offer an account in which concepts corresponding to the natural numbers are learned through a process that combines content stemming from the approximate number system with content stemming from a mechanism that tracks small numbers of objects, representing each object as being distinct from the others but without representing them as a collection—and most importantly, without representing anything numerical in relation to them. This mechanism is probably best thought of as being directed at the content domain of physical objects (or perhaps more broadly at a domain of individuals). Spelke and Tsivkin take this to be a key feature of their account, explaining why, on their account, no natural number concepts are innate—all are learned. This account is clearly a less rationalist account of the origins of natural number concepts than the one involving an innate small number system that is sketched in the text. And again, this is reflected in the degree of alignment of the resources involved with the target domain of natural numbers. As Spelke and Tsivkin's system for tracking a small number of objects is directed at the domain of objects (or individuals) and not anything numerical, it is much less closely aligned with the domain of natural numbers (though of course their account also incorporates the approximate number system). By contrast, an account involving the small number system discussed in the text is *very* closely aligned with the domain of natural numbers, since it is directed at a subset of this domain.

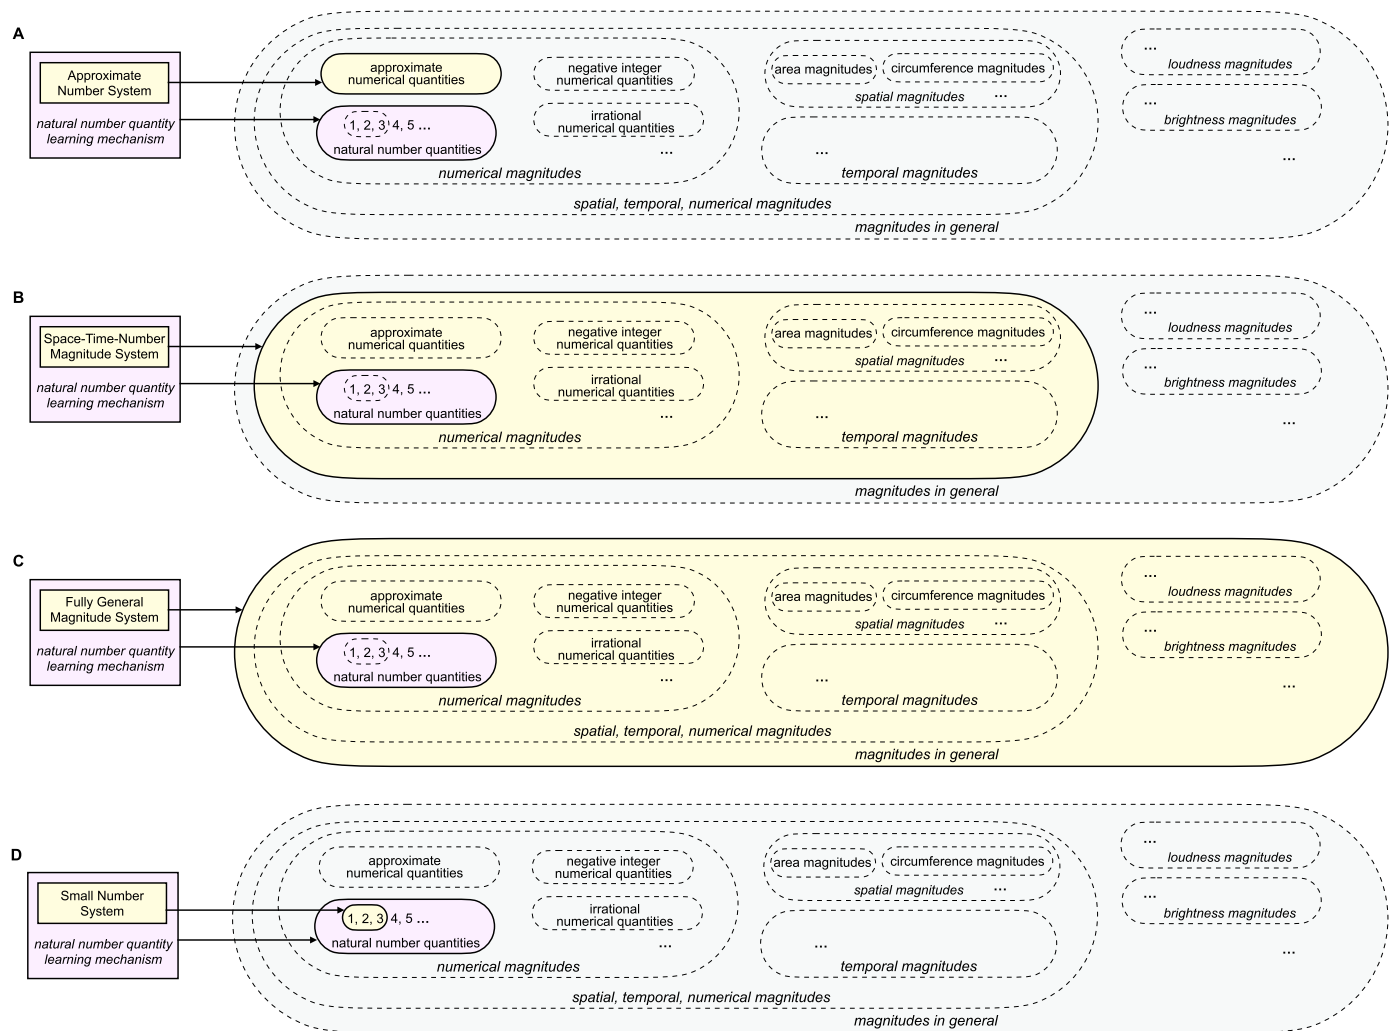


Fig. 3. Differences in alignment are correlated with how rationalist a learning mechanism is. Four schematic models of the acquisition of concepts for the natural numbers trace back to different critical innate domain-specific resources: (a) the approximate number system, (b) a system for representing spatial, temporal, and approximate numerical magnitudes in an undifferentiated manner, (c) a system for representing magnitudes-in-general, (d) a small number system, which represents the first few natural numbers as such. These postulated innate domain-specific resources are equally domain-specific, but they differ nonetheless regarding their degree of alignment to the domain targeted for learning by the learning mechanisms in question (the target domain being the domain of natural number in all four cases). The greater the degree of alignment, the more rationalist the account is.

more rationalist than another despite the fact that innate resources the two mechanisms trace back to are equally domain specific.

6. Conclusion

Domain specificity has been at the centre of some of the most important debates in cognitive science. Although numerous puzzles have arisen as a result of previous attempts to spell out exactly what this notion comes to—enough to make even advocates of domain specificity wonder if the notion is worth hanging on to—we have argued that there is no reason to abandon the notion. Domain specificity can be given a solid theoretical foundation through a greater appreciation of its perspectival nature and by getting clearer about the ways in which domain specificity interacts with related theoretical notions. This requires disentangling domain specificity from functional specificity and recognizing that domain specificity is not the only factor that contributes to the interest and significance of important domain-specific cognitive mechanisms. And in the rationalism-empiricism debate, in particular, domain specificity should be seen to work in conjunction with the degree of alignment with a target domain, among other factors. We conclude that, when seen in the proper light, domain specificity is a

perfectly coherent and invaluable theoretical construct which is fully capable of playing the major role it has been taken to have in the rationalism-empiricism debate as well as many other debates in cognitive science.

CRediT authorship contribution statement

Eric Margolis: Conceptualization, Writing—original draft, Writing—review & editing, Visualization, Funding acquisition. **Stephen Laurence:** Conceptualization, Writing—original draft, Writing—review & editing, Visualization, Funding acquisition.

Data availability

No data was used for the research described in the article.

Acknowledgements

We are grateful to the editor and to the two referees for their helpful comments and suggestions and to Canada's Social Science and Humanities Research Council and the UK's Arts and Humanities Research

Council for their support of this research.

References

- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence, & J. T. Spence (Eds.), *Vol. 2. The Psychology of Learning and Motivation: Advances in Research and Theory* (pp. 89–195). New York: Academic Press.
- Atran, S., & Medin, D. (2008). *The Native Mind and the Cultural Construction of Nature*. Cambridge, MA: MIT Press.
- Barrett, H. C. (2009). Where there is an adaptation, there is a domain: The form-function fit in information processing. In S. M. Platek, & T. K. Shackelford (Eds.), *Foundations in Evolutionary Cognitive Neuroscience* (pp. 97–116).
- Barrett, H. C. (2015). *The Shape of Thought: How Mental Adaptations Evolve*. Oxford University Press.
- Boyer, P. (2022). Ownership psychology as a cognitive adaptation: A minimalist model. *Behavioral and Brain Sciences*, 1–35.
- Carruthers, P. (2006). *The Architecture of the Mind*. Oxford University Press.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1967). Recent contributions to the theory of innate ideas. *Synthese*, 17, 2–11.
- Chomsky, N. (1972). *Language and Mind*. New York: Harcourt, Brace, Jovanovich.
- Chomsky, N. (1980). *Rules and Representations*. Columbia University Press.
- Chomsky, N. (1984). *Modular Approaches to the Study of the Mind*. San Diego State University Press.
- Cosmides, L., & Tooby, J. (1994). Origins of domain specificity: The evolution of functional organization. In L. A. Hirschfeld, & S. A. Gelman (Eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge University Press.
- Dehaene, S. (1997). *The Number Sense: How the Mind Creates Mathematics*. Oxford University Press.
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. (2000). *The Mind Doesn't Work That Way*. Cambridge, MA: MIT Press.
- Gallistel, C. R. (2000). The replacement of general-purpose learning models with adaptively specialized learning modules. In M. S. Gazzaniga (Ed.), *The Cognitive Neurosciences* (2nd ed.). Cambridge, MA: MIT Press.
- Gallistel, C. R. (2003). The principle of adaptive specialization as it applies to learning and memory. In *Principles of Learning and Memory*. Basel: Birkhäuser.
- Goodman, N. (1954). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313, 504–507.
- Hirschfeld, L. A., & Gelman, S. A. (1994a). *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge University Press.
- Hirschfeld, L. A., & Gelman, S. A. (1994b). Toward a topography of mind: An introduction to domain specificity. In L. A. Hirschfeld, & S. A. Gelman (Eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge University Press.
- Jackendoff, R. (1983). *Semantics and Cognition*. Cambridge, MA: MIT Press.
- Kamps, F. S., Julian, J. B., Battaglia, P., Landau, B., Kanwisher, N., & Dilks, D. D. (2017). Dissociating intuitive physics from intuitive psychology: Evidence from Williams syndrome. *Cognition*, 168, 146–153.
- Keil, F. C. (1981). Constraints on knowledge and cognitive development. *Psychological Review*, 88(3), 197–227.
- Keil, F. C. (1999). Nativism. In R. A. Wilson, & F. C. Keil (Eds.), *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press.
- Landau, B., & Gleitman, L. R. (1985). *Language and Experience: Evidence from the Blind Child*. Cambridge, MA: Harvard University Press.
- Laurence, S., & Margolis, E. (2005). Number and natural language. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The Innate Mind: Structure and Contents*. Oxford University Press.
- Laurence, S., & Margolis, E. (2012). The scope of the conceptual. In E. Margolis, R. Samuels, & S. Stich (Eds.), *The Oxford Handbook of Philosophy of Cognitive Science*. Oxford University Press.
- Laurence, S., & Margolis, E. (2016). Nativism, empiricism, and the interactionist consensus. *IEEE CIS Cognitive and Developmental Systems Newsletter*, 13(1), 9–10.
- Laurence, S., & Margolis, E. (forthcoming). *The Building Blocks of Thought*. Oxford University Press.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Leibovich, T., Katzin, N., Harel, M., & Henik, A. (2017). From 'sense of number' to 'sense of magnitude': The role of continuous magnitudes in numerical cognition. *Behavioral and Brain Sciences*, 40, 1–62.
- Lerner, R. M. (2015). Preface. In Lerner (Ed.), *Handbook of Child Psychology and Developmental Science* (7th edition (4 volumes)). Hoboken, NJ: Wiley.
- Leslie, A. M., & Thaiss, L. (1992). Domain specificity in conceptual development: Neuropsychological evidence from autism. *Cognition*, 43(3), 225–251.
- Lewkowicz, D. J. (2011). The biological implausibility of the nature–nurture dichotomy and what it means for the study of infancy. *Infancy*, 16(4), 331–367.
- Mahon, B. Z. (2022). Domain-specific connectivity drives the organization of object knowledge in the brain. In , *The Temporal Lobe: Vol. 187. Handbook of Clinical Neurology* (pp. 221–244). Elsevier.
- Margolis, E. (2020). The small number system. *Philosophy of Science*, 87(1), 113–134.
- Margolis, E., & Laurence, S. (2008). How to learn the natural numbers: Inductive inference and the acquisition of number concepts. *Cognition*, 106(2), 924–939.
- Margolis, E., & Laurence, S. (2013). In defense of nativism. *Philosophical Studies*, 165(2), 693–718.
- McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: Connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, 14(8), 348–356.
- Medin, D. L., & Atran, S. (1999). *Folk Biology*. Cambridge, MA: MIT Press.
- Mithen, S. J. (1996). *The Prehistory of the Mind: A Search for the Origins of Art, Religion and Science*. London: Thames and Hudson.
- Mix, K. S., Levine, S. C., & Newcombe, N. S. (2016). Development of quantitative thinking across correlated dimensions. In A. Henik (Ed.), *Continuous Issues in Numerical Cognition: How Much or How many?* Academic Press.
- Newcombe, N. S. (2002). The nativist-empiricist controversy in the context of recent research on spatial and quantitative development. *Psychological Science*, 13(5), 395–401.
- Newell, A., Shaw, J. C., & Simon, H. A. (1958). Elements of a theory of human problem solving. *Psychological Review*, 65(3), 151–166.
- Newell, A., & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Nichols, S. (2005). Innateness and moral psychology. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The Innate Mind: Structure and Contents*. Oxford University Press.
- Odic, D., & Starr, A. (2018). An introduction to the approximate number system. *Child Development Perspectives*, 12(4), 223–229.
- Prinz, J. J. (2006). Is the mind really modular? In R. Stainton (Ed.), *Contemporary Debates in Cognitive Science*. Oxford: Blackwell Publishing.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach*. MIT Press.
- Samuels, R. (2006). Is the human mind massively modular? In R. Stainton (Ed.), *Contemporary Debates in Cognitive Science*. Oxford: Blackwell Publishing.
- Santos, L., & Caramazza, A. (2002). The domain-specific hypothesis: A developmental and comparative perspective on category specific deficits. In E. M. E. Forde, & G. W. Humphreys (Eds.), *Category Specificity in Brain and Mind* (pp. 1–23). Psychology Press.
- Spelke, E. S. (2003). What makes us smart? Core knowledge and natural language. In D. Gentner, & S. Goldin-Meadow (Eds.), *Language in Mind: Advances in the Study of Language and Thought*. Cambridge, MA: MIT Press.
- Spelke, E. S., & Tsivkin, S. (2001). Language and number: A bilingual training study. *Cognition*, 78(1), 45–88.
- Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld, & S. A. Gelman (Eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge University Press.
- Sripada, C. S., & Stich, S. (2006). A framework for the psychology of norms. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The Innate Mind: Culture and Cognition*. Oxford University Press.
- Tooby, J., & Cosmides, L. (1992). The psychological foundations of culture. In J. H. Barkow, L. Cosmides, & J. Tooby (Eds.), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford University Press.
- Walsh, V. (2003). A theory of magnitude: Common cortical metrics of time, space and quantity. *Trends in Cognitive Sciences*, 7, 483–488.
- Woodward, J. F., & Cowie, F. (2004). The mind is not (just) a system of modules shaped (just) by natural selection. In C. Hitchcock (Ed.), *Contemporary Debates in Philosophy of Science*. Oxford: Blackwell Publishing.
- Yee, E., Jones, M. N., & McRae, K. (2018). Semantic memory. In J. T. Wixted, & S. Thompson-Schill (Eds.), *The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience (4th Edition, Volume 3: Language and Thought)*. New York: Wiley.