Reifying Representations

Michael Rescorla

Abstract: *The representational theory of mind* (RTM) holds that the mind is stocked with mental representations: mental items with representational properties. These items can be stored in memory, manipulated during mental activity, and combined to form complex representations. RTM is widely presupposed within cognitive science, which offers many successful theories that cite mental representations. Nevertheless, mental representations are still viewed warily in some scientific and philosophical circles. I will develop a novel version of RTM. On my approach, a mental representation is an abstract type that marks the exercise of a representational capacity. Talk about mental representations enshrines an ontologically loaded way of classifying mental states through representational capacities that the states deploy. Complex mental representations mark the appropriate joint exercise of multiple representational capacities. I will support my position with examples drawn from cognitive science, including perceptual representations and cognitive maps. I will apply my approach to longstanding debates over the existence, nature, individuation, structure, and explanatory role of mental representations.

§1. The representational theory of mind

A venerable tradition holds that the mind is stocked with *mental representations*: mental items with representational properties. There is a mental representation **whale** that represents whales, a mental representation **mammal** that represents mammals, and so on. Mental representations are similar in key respects to the communal representations employed by human society, such as pictures, maps, or natural language sentences, but they are housed in the mind rather than the external world. They can be stored in memory, manipulated during mental activity, and combined to form complex representations (e.g. the complex representation **whales are mammals**). Following Fodor (1981), let us call this picture *the representational theory of mind* (RTM).

Philosophical reception of RTM has fluctuated between acclaim and derision. In the medieval era, Ockham postulated a *language of thought* containing mental representations analogous to natural language words and sentences. Early modern philosophers frequently invoked *ideas*, often conceived in imagistic terms. In the early and mid-20th century, scientists and philosophers almost universally denounced mental representations as suspect entities that deserve no place within scientific theorizing. Tolman (1948) dissented from the anti-representationalist consensus. He hypothesized that rats navigate using *cognitive maps*: mental representations that represent the environment's spatial layout. The 1960s cognitive science revolution sparked a huge resurgence of support for mental representations, including cognitive maps (e.g. Evans, 1982; Gallistel, 1990; O'Keefe and Nadel, 1978). Fodor (1975) crystallized the trend with an influential argument that our best psychological theories presuppose a language of thought (often called *Mentalese*). Nevertheless, mental representations remain objects of suspicion in many philosophical and scientific circles.

This paper presents a novel version of RTM. I begin from the premise that the mind has *representational capacities*: a capacity to represent whales, a capacity to represent mammals, and so on. We may sort mental states and events into types based upon the representational capacities that they deploy. We may then *reify* the types, i.e. we may treat them as objects. Mental representations are the reified types. They are abstract entities that mark the exercise of representational capacities. The Mentalese word **whale** marks the exercise of a capacity to represent whales; the Mentalese word **mammal** marks the exercise of a capacity to represent mammals; and so on. Complex mental representations mark the appropriate joint exercise of representational capacities. The complex Mentalese sentence **whales are mammals** marks the appropriate joint exercise of the representational capacity corresponding to **whale**, the representational capacity corresponding to **mammal**, and other capacities as well. A cognitive map marks the appropriate joint exercise of capacities to represent individual landmarks, a capacity to represent oneself, and a capacity to represent entities as spatially positioned a certain way. I call my approach *the capacities-based representational theory of mind* (*C-RTM*).

I will develop C-RTM and apply it to longstanding debates over the existence, nature, structure, individuation, and explanatory role of mental representations. §2 motivates a foundational commitment underlying my treatment: mental states and events have explanatorily significant representational properties. §3 presents my core thesis: mental representations are abstract types that we cite so as to classify mental states and events in representational terms. Anyone who accepts that mental states can represent the world and who countenances abstract entities should countenance mental representations as construed by C-RTM --- or so §4 argues. §5 uses C-RTM to elucidate the familiar thought that some mental representations are *complex*. §§6-7 address how to individuate mental representations, with particular emphasis upon the individuative role that C-RTM assigns to representational properties.

§2. Mental representation and mental representation*s*

Researchers across philosophy and cognitive science use the phrase "mental representation" in many different ways. On the usage that concerns me, mental representation is tied to *veridicality-conditions*: conditions for veridical representation of the world. To illustrate:

- Beliefs are the sorts of things that can be true or false. My belief *that Emmanuel Macron is French* is true if Emmanuel Macron is French, false if he is not.
- Perceptual states are the sorts of things that can be accurate or inaccurate. My perceptual experience *as of a red sphere before me* is accurate only if a red sphere is before me.
- Desires are the sorts of things that can be fulfilled or thwarted. My desire *to eat chocolate* is fulfilled if I eat chocolate, thwarted if I do not eat chocolate.

Beliefs have truth-conditions, perceptual states have accuracy-conditions, and desires have fulfillment-conditions. Truth, accuracy, and fulfillment are species of veridicality. So beliefs, perceptual states, and desires have veridicality-conditions.

In daily life, we often explain mental and behavioral outcomes by citing beliefs, desires, and other representational mental states. We identify these mental states through their veridicality-conditions or through representational properties that contribute to veridicality-conditions. When we say "Frank believes that Emmanuel Macron is French," we specify the condition under which Frank's belief is true (that Emmanuel Macron is French). When we say "Frank wants to eat chocolate," we specify the condition under which Frank's desire is fulfilled (that Frank eats chocolate). Everyday discourse assigns a central role to *intentional explanations*, i.e. explanations that cite veridicality-conditions or representational properties that contribute to veridicality-conditions. We can formalize many aspects of folk psychological intentional explanation more rigorously using the mathematical apparatus of Bayesian decision theory, which replaces the binary notion *belief* with the graded notion *degree of belief* and the binary notion *desire* with the graded notion *utility*.

Contemporary science builds upon folk psychology by assigning mental representation a central role within psychological explanation. A few examples:

- *High-level cognition*, including belief-fixation, decision-making, deductive reasoning, planning, problem solving, linguistic communication, and so on. Folk psychology routinely explains high-level mental phenomena by citing representational mental states, such as beliefs, desires, and intentions. The same explanatory strategy figures prominently within social psychology, developmental psychology, economics, linguistics, and all other fields that study high-level cognition. Each field presupposes mental states resembling those posited by folk psychology. Each field takes folk psychology (sometimes filtered through Bayesian decision theory) as a template.
- *Perception*. Cognitive science extends folk psychology by describing *subpersonal* processes in representational terms. For example, *perceptual psychology* studies how the perceptual system transits from proximal sensory stimulations to perceptual states that represent environmental conditions. The orthodox view, going back to Helmholtz (1867), holds that the transition involves an "unconscious inference" from proximal stimulations to representational perceptual states. The inference is executed not by the person herself but rather by her mental subsystems. To describe unconscious perceptual inference, contemporary perceptual psychologists offer mathematical models grounded in Bayesian decision theory (Knill and Richards, 1996). The models cite representational properties of perceptual states, such as representational relations to shapes, sizes, colors, and other aspects of the distal environment (Rescorla, 2015).

- *Navigation*. Cognitive science extends the intentional paradigm from humans to non-human animals. Psychologists attribute representational mental states to mammals, many birds, and some insects. For example, the research program launched by Tolman (1948) has established that spatial representation underwrites mammalian navigation. Mammals represent how the environment is spatially arranged, and on that basis they navigate through space (Rescorla, 2018).

Overall, intentional explanation has achieved striking success within cognitive science. It illuminates *perception* (Burge, 2010a), *motor control* (Rescorla, 2016), *navigation* (Rescorla, 2018), *deductive reasoning* (Rips, 1994), *mathematical cognition* (Gallistel and Gelman, 2005), *language* (Heim and Kratzer, 1998), and many other core mental phenomena. Bayesian models of the mind have proved especially successful (Rescorla, 2020b; Rescorla, forthcoming a).

This scientific work provides strong abductive support for *intentional realism*: realism about representational mental states and events. Intentional realists hold that representational properties are genuine, scientifically important features of mentality. Fodor defends intentional realism in a series of publications stretching over several decades (1975, 1981, 1987, 1994, 2008). Burge (2010a) also defends intentional realism, focusing on perception.

Over the past century, many philosophers and scientists have opposed intentional realism (e.g. Chemero, 2009; Churchland, 1981; Quine, 1960; Skinner, 1938; Stich, 1983; van Gelder, 1992). Anti-representationalists hold that we should eschew talk about mental representation, just as modern chemistry eschews talk about phlogiston. I agree with Fodor and Burge that the anti-representationalist perspective is misguided. Philosophical arguments against intentional explanation are uniformly unconvincing. Attempts at expunging representationality from science have failed, with the proposed anti-representationalist theories typically far less successful than the representationalist theories they are meant to supplant. Overall, current cognitive science offers numerous impressive intentional explanations whose benefits do not look replicable within a non-representational approach. I will not engage any further with anti-representationalism. I focus instead on developing intentional realism.

Intentional realism does not entail RTM. One can acknowledge the reality and importance of representational mental states while rejecting all talk about mental representation*s*. Soames (2010) and Stalnaker (1984) occupy this position. When I perspire, we do not postulate perspiration*s* that do the perspiring. When I procreate, we do not postulate procreation*s* that do the procreating. When I relax, we do not postulate relaxation*s* that do the relaxing. Why, then, should mental representation involve mental representation*s* that do the representing? The usual rejoinder is to defend RTM abductively. Within psychology, Tolman (1948) argues that cognitive maps help us explain mammalian navigation. Within philosophy, Fodor (1975) contends that our best cognitive science explanations presuppose a language of thought.

I think that abductive arguments for RTM have considerable force. However, I find them most compelling when they specify as carefully as possible what is being claimed when one postulates mental representations. In particular, a truly satisfying abductive argument should elucidate how exactly RTM goes beyond intentional realism. §§3-4 offer my own preferred elucidation.

§3. Mental representations as types

I will pursue the following intuitive idea: mental representations are types that mark the exercise of representational capacities.

In quotidian and scientific discourse, we incessantly classify entities into categories. We taxonomize sounds, shapes, utterances, actions, artifacts, molecules, animals, books, symphonies, dances, and so. Given any reasonable scheme for taxonomizing entities, we may recognize a collection of *types* corresponding to the taxonomic scheme. Each type correlates with a category employed by the taxonomic scheme. We *type-identify* an entity by specifying a type instantiated by the entity. Types are *abstract objects*. When I say that they are "objects," I mean that we can refer to them, quantify over them, and count them (Parsons, 2008). When I say that they are "abstract," I mean that they are not located in space and time and that they do not participate in causal interactions. An important feature that distinguishes types from other abstract entities --- such as sets, numbers, and functions --- is that types are instantiated by *tokens*. Individual copies of *Moby Dick* are tokens of a single book-type. Individual dogs are tokens of a single species. Individual rectangles are tokens of a single shape. Suitable inscriptions or utterances are tokens of the English word-type "dog." In each case, we posit an abstract type instantiated by certain tokens. We *reify* types by treating them as objects.[1]

Why reify? Why posit abstract types instantiated by tokens? One reason is that reification of types pervades virtually all serious discourse. As Wetzel (2009) documents, types figure pervasively across a vast range of human endeavors. A very partial list: linguistics (*words*, *sentences*, *phonemes*); computer science (*LISP expressions*, *computational systems*); biology (*genes*, *species*); chemistry (*atoms*, *molecules*); physics (*forces*, *fields*, *protons*); music theory (*notes*, *chords*, *arpeggios*); politics (*bills*, *parliamentary procedures*); athletics (*gymnastics routines*, *football plays*); gaming (*chess gambits*, *poker hands*). In most cases, it is doubtful that we could preserve anything like current expressive power while expunging types from our

[1] On reification in general, see Quine's writings (e.g., 1980, 1981, 1995).

discourse. For example, each of the following statements refers to, quantifies over, or counts types:

> Chaucer introduced fewer words into the English language than Shakespeare.
>
> Over 1,000 species went extinct last year.
>
> Mary learned numerous chess gambits yesterday, including the queen's gambit and the king's gambit.
>
> The gas in the container is composed of five different molecules, including methane and carbon dioxide.
>
> During the music theory lecture, Professor Smith discussed the diminished seventh chord, the Neapolitan sixth chord, and four other new chords.
>
> The filibuster is arguably the most pernicious parliamentary procedure regularly used in the United States Senate.

In each case, it is unclear how one might paraphrase away the reified types. Even if it is possible in principle to paraphrase types away, doing so would as a practical matter be an oppressive burden. Taking current practice as our guide, types seem as useful and firmly entrenched within our discourse as any other entities, including material objects.

Let us apply this viewpoint to the special case of mental representation. As I argued in §2, cognitive science frequently classifies mental events in representational terms. Just as other disciplines reify types, so can cognitive science reify mental event types. Given a taxonomic scheme that classifies mental events through their representational properties, we can posit a corresponding collection of types. Each type correlates with a category employed by our representationally-based taxonomic scheme. On my version of RTM, mental representations are the reified representational mental event types. Their tokens are token mental events. Here I construe the term "event" broadly to include states, such as beliefs, and processes, such as inferential transitions within thought.

I will develop these ideas by invoking the notion of a *representational capacity*. Whenever we describe a creature's mental activity in representational terms, we presuppose that the creature has certain representational capacities. Examples:

- *High-level cognition*. When a thinker forms the belief *that some dogs are furry*, he exercises a capacity to represent dogs. By describing his mental state as a belief *that some dogs are furry*, we recognize (at least implicitly) that he has exercised this capacity. When we describe him as forming an intention *to eat chocolate*, we recognize (at least implicitly) that he has exercised a capacity to represent chocolate.
- *Perception*. Perceptual states represent shapes, sizes, colors, and other distal properties. By describing perceptual states in this way, we presuppose perceptual capacities to represent shapes, sizes, colors, and so on. For example, by describing a perceptual state as an estimate that an object is spherical, we recognize (at least implicitly) that the perceiver has exercised a perceptual capacity to represent sphericality.
- *Navigation*. Mammals mentally represent the environment's spatial layout, especially the locations of salient landmarks. By saying that a creature mentally represents a specific spatial layout, we recognize (at least implicitly) that the creature has exercised a capacity to represent that layout.

Cognitive science routinely groups a creature's mental events into types by invoking (at least implicitly) representational capacities exercised by the creature or the creature's mental subsystems. We may posit an array of types corresponding to those representational capacities:

- *High-level cognition*. Suppose a thinker has a capacity to represent dogs within high-level thought. The thinker might exercise this capacity by forming a belief *that some dogs are furry*, or a desire *to buy a new dog*, or a conjecture that *dogs are closely related to wolves*, and so on. We may classify the thinker's mental events based upon whether she exercises the capacity. We may then reify, positing a type **dog** instantiated precisely when the thinker exercises the capacity. Citing the type **dog**, we type-identify mental events based upon whether they deploy the correlated representational capacity. In similar fashion, we may recognize a mental representation **whale**, a mental representation **furry**, and so on. These items resemble the *Mentalese words* posited by Fodor (although §7 will argue that they do not have all the properties attributed by Fodor to Mentalese words).
- *Perception*. Suppose a perceiver has a capacity to represent sphericality within perception. This capacity can figure in innumerable perceptual states. For example, the perceiver might perceptually represent an object as a large red sphere, or she might perceptually represent a different object as a small green sphere, and so on. We may classify perceptual states based upon whether the perceiver exercises the capacity. We may then reify by recognizing a *perceptual representation* instantiated precisely when the perceiver exercises the capacity. Similarly for other perceptual capacities, such as capacities to represent specific sizes, colors, etc.
- *Navigation*. Suppose an animal has capacities to represent a range of possible spatial layouts. We may classify the animal's mental events based upon which such capacity the animal exercises. We may posit a range of mental representations, each

instantiated precisely when the animal exercises the correlated capacity. These mental representations resemble the *cognitive maps* posited by Tolman.

According to *the capacities-based representational theory of mind* (*C-RTM*), a mental representation is a type that correlates with a representational capacity. The type is instantiated precisely when a thinker (or one of her subsystems) exercises the correlated capacity. As I will put it, the type *marks* the exercise of the representational capacity.[2] C-RTM asserts that mental representations as thus construed exist and are explanatorily important.

From C-RTM's viewpoint, talk about mental representations embodies an "ontologized" way of classifying mental events through representational capacities deployed by the events. Such talk is a more ontologically loaded expression of something to which folk psychology and cognitive science are firmly committed: animals have representational capacities, and those capacities are important for understanding how the mind works. By positing mental representations, we reify the mental event types that figure in our quotidian and scientific classificatory procedures. We thereby undertake ontological commitment to the types. In some areas, such as research on cognitive maps, current theorizing already makes the reification explicit. In other areas, such as folk psychology, explicit reification is less common.

When introducing C-RTM, I used the locution "exercise a representational capacity." A few clarificatory remarks about this locution and its role in my account:

- I do not construe the locution in a voluntaristic or action-theoretic way. You can exercise a capacity involuntarily, without intending to exercise it, and without being consciously aware that you are exercising it. For example, you have a capacity to breathe, and while reading this paper you have until now most likely been exercising that capacity involuntarily, without intending to, and without conscious awareness.

[2] I borrow the term "mark" from Burge (2010a, p. 39), who uses it in a similar but not completely identical fashion.

Many of the representational capacities invoked by cognitive science are not under voluntary control: you cannot help but perceive a sphere as spherical when you see it under the proper viewing conditions. In some cases (such as Bayesian inferences executed by the perceptual system), a representational capacity is exercised not by the thinker but rather by a mental subsystem. Exercising a representational capacity is not necessarily something that you decide to do, something you know you are doing, something you can control, or even something done by you (as opposed to your mental subsystems).

- You may *have* a representational capacity without *exercising* it. For example, many people have an unexercised capacity to represent the property *green-eyed ophthalmologist currently living in Marseille who loves rock-climbing*. We may posit a mental representation **green-eyed ophthalmologist currently living in Marseille who loves rock-climbing** that marks the exercise of this representational capacity. Someone who has the capacity without exercising it never instantiates the mental representation. Had she exercised the capacity, she would have instantiated the representation. A mental representation is an abstract type that exists and that marks the exercise of a representational capacity *whether or not the capacity is actually exercised*.
- When a thinker (or one of her subsystems) exercises a representational capacity, there occurs a mental event that is an exercise of the representational capacity. The event might be a judgment, a belief, a perceptual state, an inferential transition, and so on. A mental event is a token of a mental representation just in case the event is an exercise of the representational capacity marked by the representation.

It would be good to say more about capacities in general, about representational capacities more specifically, and about the relation *mental event e is an exercise of representational capacity C* even more specifically. My strategy in the present paper is to take these notions for granted and explore whether they can help clarify the nature of mental representations.

It does not immediately follow from what I have said that mental representations have representational properties. A type does not automatically inherit properties from its tokens. For example, the type *red ball* has red tokens but is not itself red. Likewise, one might hold that mental representation *R* marks the exercise of a capacity to represent *d* but that *R* does not itself represent *d*. I finesse this issue by introducing a term *denote** governed by the following stipulation:

Mental representation *R* denotes* *d* iff *R* marks the exercise of a capacity to represent *d*.

There seems little harm in writing the "*" with invisible ink, yielding:

(Δ) Mental representation *R* denotes *d* iff *R* marks the exercise of a capacity to represent *d*.

We may therefore postulate mental representations that denote, keeping in mind that the operative notion of "denotation" is given by (Δ).

§4. Admitting mental representations into our discourse

C-RTM goes beyond intentional realism in two main ways. First, C-RTM invokes representational capacities. Strictly speaking, an intentional realist might eschew all talk about representational capacities. However, this first difference does not strike me as very significant, because invocation of representational capacities seems at least implicit in any view that attributes representational properties to mental events. Second, and more importantly, C-RTM outstrips intentional realism by explicitly reifying representational mental event types. A theorist might type-identify mental events representationally while declining to reify. For example, any intentional realist should agree that

a belief *that some dogs are furry*.

a desire *to buy a new dog*.

a conjecture *that dogs are closely related to wolves*.

a fear *that Fred's dog will bark.*

etc.

have something important in common. A common capacity to represent dogs is exercised in each case. C-RTM goes further by asserting that *there exists* a type **dog** instantiated in each case. The existential quantifier signals ontological commitment to an abstract type. By reifying representational mental event types, C-RTM goes beyond intentional realism.

Why reify mental representations? Why posit representational mental event types?

Well, why not? We usually feel no qualms about positing abstract types corresponding to a reasonable taxonomic scheme. Why feel qualms about reifying the mental event types that figure in intentional explanation? (Cf. Burge, 2010a, p. 40.) Given the plethora of types recognized across a vast range of human endeavors, and given the scientific success of representational taxonomization, why scruple at positing representational mental event types?

Philosophers who favor a broadly *nominalist* viewpoint will decry the reifying step from intentional realism to C-RTM. Nominalists deny that abstract entities exist. In particular, they deny that types exist. However, nominalism is a problematic position. A large literature over the past century has convincingly demonstrated that abstract entities are, at least in some cases, metaphysically harmless (Linnebo, 2018; Parsons, 2008) and indispensable to scientific inquiry (Quine, 1981, pp. 1-23, pp. 148-155; Putnam, 1975a, pp. 345-357). Setting nominalist worries aside, I see no good reason why we should decline to reify representational mental event types. Any theorist who is comfortable with abstract entities in general should gladly take the extra reifying step from intentional realism to C-RTM.

I distinguish two dialectical roles for abduction within the defense of C-RTM. First, as I urged in §2, there is strong abductive evidence for intentional realism. Having endorsed intentional realism, we should happily take the extra reifying step to C-RTM. Second, there are several areas where our best cognitive science theories already take the reifying step by assigning explanatory centrality to mental representations. Examples:

- According to Bayesian perceptual psychology, the perceptual system attaches probabilities to *hypotheses*. For instance, shape perception results from a Bayesian inference over hypotheses that represent possible distal shapes. Bayesian models individuate hypotheses representationally --- by citing shapes, sizes, colors, and other such distal properties represented by the perceptual system (Rescorla, 2015). A Bayesian model of shape perception individuates hypotheses in terms of represented shapes; a Bayesian model of size perception individuates hypotheses in terms of represented sizes; and so on. The science presupposes standing capacities for perceptual representation of distal properties. It invokes these standing capacities when individuating hypotheses that figure in Bayesian perceptual inference. Thus, Bayesian perceptual psychology posits perceptual representations as characterized by C-RTM.
- Researchers posit cognitive maps employed during mammalian navigation. Detailed computational models, some Bayesian, describe how mammals use sensory input and self-motion cues to update their cognitive maps and how they use cognitive maps

> during navigation (Madl, et al. 2015). The models individuate cognitive maps representationally --- by citing specific represented spatial layouts (Rescorla, 2018). The science presupposes that mammals have standing capacities to represent spatial layouts. It invokes these capacities when individuating cognitive maps. Thus, current research posits cognitive maps as construed by C-RTM.

Explicit reification of representational mental event types also occurs in explanatorily fruitful theories of motor control (Rescorla, 2016), deductive reasoning (Rips, 1994), mathematical cognition (Gallistel and Gelman, 2005), and many other phenomena.

In most such cases, it is unclear whether we can preserve the theory's explanatory achievements while declining to reify. Take Bayesian perceptual psychology. Bayesian perceptual models postulate hypotheses to which probabilities attach. To accept that the models are true or even just approximately true, one must accept that the postulated hypotheses exist and are roughly as depicted by the models. The models individuate hypotheses representationally, by invoking representational capacities exercised when the hypotheses are instantiated. A non-representational individuative scheme would depart dramatically from current science, most likely with undesirable explanatory repercussions (Rescorla, 2015). When a successful theory reifies representational types, and when rejecting the types would beget an explanatory loss, this provides abductive evidence for C-RTM.

C-RTM is a very anodyne version of RTM, in that it involves fairly minimal commitments beyond intentional realism. Some readers may fear that it is *too* anodyne. How much substance can there be to a version of RTM that extends so little beyond intentional realism? Can such a weak view really capture what proponents of RTM have intended by postulating mental representations, or what opponents have intended by denying that such entities exist? My appeal to representational capacities may also appear suspiciously empty. When I say that mental event *e* is an exercise of a capacity to represent *d*, isn't that just a needlessly prolix way of saying that *e* represents *d*? It may seem that representational capacities are doing no real work, so that my account has even less substance than I have advertised.

The rest of the paper addresses these worries. My basic strategy is to show that the appeal to representational capacities advances our understanding of mental representations, sometimes in surprising ways. §5 uses C-RTM to clarify the *complexity* of certain mental representations. §6 uses C-RTM to elucidate how distinct mental representations can share the same denotation. §7 argues that C-RTM improves upon many rival versions of RTM by assigning a central individuative role to representational properties. §§5-7 collectively show that C-RTM, properly developed, is a highly substantive view that preserves and illuminates many traditional core commitments of RTM.

§5. Complex mental representations as complex types

Advocates of RTM universally agree that mental representations can combine to form complex representations. What does it mean to say that a mental representation is "complex"? I think that C-RTM supplies a good answer. Before saying how, I must reflect in a general way upon types and tokens.

§5.1 Complex types

It is often natural to regard a type as a complex entity that bears structural relations to other types. The molecule methane CH_4 is instantiated when four hydrogen atoms and one carbon atom form appropriate chemical bonds with one another. A C-major chord occurs when the three notes C, E, and G play simultaneously. The English sentence "John loves Mary" occurs when the individual words "John," "loves," and "Mary" are arranged in a suitable syntactic structure. In each case, we recognize a type that marks the structured instantiation of other types. The first type intimately involves the other types as arranged in an appropriate configuration. What counts as "appropriate" varies. The same notes (C, E, and G) yield a chord or an arpeggio depending on whether they play simultaneously or sequentially.

I say that type *t* *incorporates* types $t_1, \ldots, t_n$ other than itself when, necessarily, any token of *t* has parts that are tokens of $t_1, \ldots, t_n$ and that bear appropriate relations to one another. The "appropriate relations" depend upon *t*. I say that type *t* is *complex* when it incorporates other types. Here I employ an abstract notion of *part* affiliated with mereology. Parts need not be spatial parts. For example, *the Biles* is a complex gymnastics routine first performed by Simone Biles consisting of a double layout and half twist, where the half twist occurs during the end of the double layout. Each token of the Biles is a complex event with two distinct parts: a token performance of a double layout, and a token performance of a half twist. A token performance of a double layout is part of the overall token performance of the Biles, but it is not a *spatial* part. Recognizing a type as complex is often an essential first step towards elucidating it. That is why, when we introduce a novice to a complex type such as a C-major chord, methane, or the Biles, we usually cite other types and indicate how their tokens must relate in order for the complex type to be instantiated.

A complex type has tokens that are themselves complex. The tokens may be medium-sized physical objects (e.g. automobiles), microscopic particles (e.g. molecules), events (e.g. linguistic utterances, gymnastic performances), or otherwise. Each token of a complex type *t* has parts that are token of the types incorporated by *t*.

Complex types resemble the *structural universals* posited by Armstrong (1980). Lewis helpfully characterizes structural universals as follows (1999, p. 81): "Anything that instantiates [a structural universal] must have parts; and there is a necessary connection between the instantiating of the structural universal by the whole and the instantiating of other universals by the parts." However, I do not assume that complex types have all the properties attributed by Armstrong and Lewis to universals. More specifically, Lewis (1999, p. 80) says that a universal satisfies two conditions: "wherever it is instantiated, there the whole of it is present" and "[w]hen it is instantiated, it is a nonspatiotemporal part of the particular that instantiates it." I do not assume that types satisfy either condition.[3]

Does a complex type have parts other than itself? Does it have internal structure? It sounds plausible to say that methane is composed of carbon and hydrogen bound together in an appropriate way, that a natural language sentence is composed of certain words arranged in a certain syntactic structure, and so on. Such claims are common within philosophy, linguistics, computer science, and most other disciplines that cite complex types. However, they do not follow from my definition of "complex type." My definition requires that *tokens* of the complex type have parts, not that the *complex type itself* have parts. One might hold that a complex type is an unstructured entity whose tokens necessarily have internal structure. One might hold that complex type t marks the structured instantiation of $t_1, \ldots, t_n$ even though t is not itself structured. I doubt that any truly important questions about the mind hinge upon how we resolve this issue. What matters for my purposes is the role that complex types play within our discourse: a

[3] Lewis (1999, pp. 78-107) argues that structural universals do not exist. His arguments generalize straightforwardly from structural universals to complex types. This is worrisome, because complex types figure across a vast range of disciplines, including linguistics, computer science, biology, chemistry, music theory, and so on. Were Lewis's arguments successful, we would face the unsavory prospect of revising these disciplines so as to eliminate reference to complex types. Luckily, the literature suggests several ways that one might contest Lewis's arguments (e.g. Bennett, 2013; Davis, 2014; Hawley, 2010). For present purposes, we may safely disregard Lewis's arguments and assume that complex types exist.

complex type marks the structured instantiation of other types. Whether the complex type itself has structure is an interesting question that I set aside.[4]

A complex type is instantiated only when incorporated types are instantiated in appropriate relation to one another. In that sense, the complex type is instantiated only when incorporated types combine in an appropriate way. While I remain agnostic as to whether complex types have internal structure, it still seems right to say that a complex type results from combining types that it incorporates. We can say with good conscience that the C-major chord results from appropriately combining C, E, and G, that the Biles results from appropriately combining a double-layout and a half-twist, and so on. In agreeing that a complex type results from combining incorporated types, we need not agree that incorporated types are *parts* of the complex type. There are many cases where an entity results from combining items that are not its parts. A cake results from appropriately combining its ingredients, but few philosophers would say that the ingredients are parts of the cake.

§5.2 Structured instantiation of representational capacities

Typically, a representational mental event deploys multiple representational capacities in coordination with one another. The coordinated capacities are sub-capacities of a *complex representational capacity*. A *complex mental representation* is a mental representation that marks the exercise of a complex representational capacity. To illustrate:

- *High-level cognition*. When a thinker judges *that some dogs are furry and some are not*, she exercises a complex representational capacity composed of sub-capacities that include a capacity to represent dogs (corresponding to the Mentalese word **dog**),

[4] The two possible answers to this question correspond to two possible views distinguished by Lewis (1999) regarding structural universals: *the pictorial conception* ("a structural universal is isomorphic to its instances" – p. 96) and *the magical conception* ("a structural universal has no proper parts" – p. 100).

a capacity to represent furriness (corresponding to the Mentalese word **furry**), and capacities for conjunction, negation, and existential quantification. We posit a complex Mentalese sentence **some dogs are furry and some are not**, which marks the exercise of the complex representational capacity.

- *Perception*. A typical perceptual state attributes distal properties to observed particulars (Burge, 2010b). The perceptual state deploys a complex representational capacity composed of sub-capacities that may include: capacities to represent specific shapes, sizes, colors, and other distal properties; capacities for singular representation of environmental particulars; and a capacity to attribute distal properties to perceived particulars (Rescorla, 2020a). We may posit a complex perceptual representation that marks the exercise of this complex representational capacity.
- *Navigation*. A cognitive map marks the exercise of a complex representational capacity composed of sub-capacities that include: capacities to represent individual landmarks; a capacity to represent the animal itself; and capacities to represent particulars as positioned a certain way in physical space.

A complex mental representation marks the coordinated exercise of representational capacities.

I now argue that complex mental representations are complex types. A complex mental representation *R* marks the exercise of a complex representational capacity *C*. A mental event is a token of *R* just in case it is an exercise of *C*:

(1) Necessarily, any token of *R* is an exercise of *C*.

C has sub-capacities C_1, …, C_n. Exercise of *C* consists in the coordinated exercise of C_1, …, C_n. Exercise of C_i is part of the exercise of *C*, just as a gymnast's performance of a double-layout is part of her performance of the Biles. Any exercise of C has as parts exercises of sub-capacities $C_1, \ldots, C_n$, where the sub-capacities must be coordinated appropriately with one another:

(2) Necessarily, any exercise of C has parts that are exercises of $C_1, \ldots, C_n$ and that bear appropriate relations to one another.

There exist mental representations $R_1, \ldots, R_n$ that mark the exercise of capacities $C_1, \ldots, C_n$:

(3) Necessarily, any exercise of C_i is a token of R_i.

(1)-(3) entail:

(4) Necessarily, any token of R has parts that are tokens of $R_1, \ldots, R_n$ and that bear appropriate relations to one another.

Since $R_1, \ldots, R_n$ are distinct from R, it follows from (4) that R incorporates $R_1, \ldots, R_n$. Thus, any complex mental representation R is a complex type in §5.1's sense. R's tokens are complex mental events.[5]

Clause (4) is a crucial feature of complex mental representations, found across a wide range of psychological domains. Examples:

- *Mentalese sentences*. To instantiate **some dogs are furry and some are not**, a thinker must instantiate the type **dog**, the type **furry**, and types corresponding to conjunction, negation, and the existential quantifier. The thinker must instantiate these types in appropriate relation to one another. So the Mentalese sentence incorporates **dog**, **furry**, and types corresponding to conjunction, negation, and the existential quantifier.
- *Complex perceptual representations*. Suppose a token perceptual state attributes observable properties to observed particulars. Then the state instantiates a complex

[5] Fodor also explicates the complexity of mental representations by citing the complexity of mental events (1987, pp. 136-139). However, he combines his emphasis on complex mental events with several additional doctrines that I do not accept (see §5.5 and §7 below).

perceptual representation that incorporates perceptual representations of the properties and singular perceptual representations of the particulars.

- *Cognitive maps*. To instantiate a cognitive map, an animal must instantiate various *singular representations* that mark the exercise of capacities to represent various particulars, including both individual landmarks and the animal itself. The animal must also instantiate *mental coordinates* that mark the exercise of capacities to represent physical particulars as positioned a certain way in physical space (Gallistel, 1999). The animal instantiates these types in appropriate relation to one another, thereby representing the denoted particulars as positioned a certain way in space.

In each case, a complex mental representation R marks the structured instantiation of mental representations $R_1, \ldots, R_n$.

§5.3 A toy example

Let me illustrate these ideas with a toy example. Imagine an idealized mathematical thinker with a capacity to represent the number 0 and a capacity to represent the successor function. I posit mental words 𝟘 and 𝕊 that mark the respective exercise of each capacity. The thinker also has a capacity to apply functions to arguments. These three capacities yield complex capacities to represent natural numbers. I posit an array of complex mental numerals that mark the exercise of the complex capacities. For example, there is a mental numeral 𝕊𝕊𝕊𝕊𝟘 that marks the exercise of a complex capacity to represent the number 4. This complex capacity involves three sub-capacities:

a capacity to represent 0

a capacity to represent the successor function (deployed four times) a capacity to apply a function to an argument (deployed four times)

To instantiate 𝕊𝕊𝕊𝕊𝟘, a thinker must instantiate 𝟘 and 𝕊, and she must do so in the appropriate way --- by iteratively combining her capacity for function-application with the capacities corresponding to 𝟘 and 𝕊. Thus, 𝕊𝕊𝕊𝕊𝟘 incorporates 𝟘 and 𝕊.

Similarly, suppose the thinker has a capacity to represent the addition function. I posit a mental word **+** correlated with this capacity, and I posit further complex mental numerals that mark the exercise of the resulting complex capacities. Mental numeral 𝕊𝕊𝟘 **+** 𝕊𝕊𝟘 marks the exercise of a complex capacity involving four sub-capacities:

a capacity to represent 0

a capacity to represent the successor function (deployed four times)

a capacity to represent the addition function (deployed once)

a capacity to apply a function to an argument (deployed five times)

This complex capacity is a capacity to represent the number four. However, it is a different capacity than the complex capacity corresponding to 𝕊𝕊𝕊𝕊𝟘. The capacities are different because they involve different sub-capacities exercised in different ways.

More generally, the thinker has an infinite array of representational capacities, corresponding to an infinite array of mental numerals. If *t* is mental numeral, then 𝕊*t* is a complex mental numeral that marks appropriate joint exercise of the capacity corresponding to 𝕊, the capacity corresponding to *t*, and a capacity for function-application:

(5) If mental numeral *t* marks the exercise of a capacity to represent *d*, then 𝕊*t* marks the exercise of a capacity to represent the successor of *d*.

If *s* and *t* are mental numerals, then *s* **+** *t* is a complex mental numeral that marks appropriate joint exercise of the capacity corresponding to **+**, the capacities corresponding respectively to *s* and *t*, and a capacity for function-application:

(6) If mental numeral *s* marks the exercise of a capacity to represent *d*, and if mental numeral *t* marks the exercise of a capacity to represent *e*, then *s* **+** *t* marks the exercise of a capacity to represent the sum of *d* and *e*.

What results is a toy language of complex mental words, each word correlated with a distinct complex representational capacity.

Our toy language comes with a canonical compositional semantics. We stipulated that 𝟘 marks the exercise of a capacity to represent 0. Recall also our stipulation from §3:

(Δ) Mental representation *R* denotes *d* iff *R* marks the exercise of a capacity to represent *d*.

Our stipulations jointly entail the base clause:

𝟘 denotes 0.

(5) and (Δ) jointly entail the recursion clause:

𝕊*t* denotes the successor of the denotation of *t*,

for any mental numeral *t*. (6) and (Δ) jointly entail the recursion clause:

*s***+***t* denotes the sum of the denotation of *s* and the denotation of *t*,

for any mental numerals *s* and *t*. From the base clause and recursion clauses, we can derive familiar Tarski-style clauses specifying the denotations of specific terms, such as:

𝕊𝕊𝕊𝕊𝟘 denotes the successor of the successor of the successor of the successor of 0.

Our stipulations thereby determine a unique denotation for each mental numeral, as specified by the compositional semantics.

§5.4 Complex mental representations

I propose that we take the toy mathematical language as a paradigm. When studying a mental phenomenon --- perception, motor control, navigation, language, high-level cognition, or what have you --- we should identify the complex representational capacities at work. We should clarify how each complex capacity decomposes into sub-capacities. We may then posit complex mental representations corresponding to the complex capacities. For each complex mental representation, we should clarify which sub-capacities are exercised when the representation is instantiated, which "appropriate relations" must obtain between the exercised sub-capacities, and which representational properties are implicated by appropriately coordinated exercise of the sub-capacities.

Because the toy mathematical language is an artificial example, I could stipulate its properties. When studying real-life examples, we may no longer stipulate. We must instead seek guidance from folk psychology, scientific psychology, introspection, and philosophical reflection. In some cases, those sources already afford considerable insight. Two examples:

- Any thinker has a general capacity for *predication*. If she has a capacity to represent n-place relation F, and if she has a capacity to represent objects $a_1, \ldots, a_n$, then her general predicative capacity also makes available (in principle) a complex capacity for predicating F of $a_1, \ldots, a_n$. Exercising this capacity, she can think a thought that is true just in case $a_1, \ldots, a_n$ stand in relation F to one another. For example, given a capacity to represent London, a capacity to represent Paris, and a capacity to represent the relation in which two entities stand when the first is north of the second, the agent can think a thought that is true just in case London is north of Paris. In doing so, she exercises a complex capacity composed of the aforementioned sub-capacities. Exercise of the complex capacity is marked by a Mentalese sentence **London is north of Paris** that incorporates individual Mentalese expressions **London**, **Paris**, and **is north of**. In order for the sentence to be instantiated, the individual expressions must be instantiated in the appropriate way, drawing upon the thinker's predicative capacity.

- Any thinker has a general capacity for *conjunctive thought*. Given a capacity to think a thought with some truth-condition, and given a capacity to think a thought with a second truth-condition, the thinker has a capacity to think a thought that is true iff both truth-conditions are satisfied. Accordingly, we may posit a Mentalese word **and** that combines appropriately with Mentalese sentences. The Mentalese sentence **London is north of Paris and Paris is north of New York** correlates with a complex capacity composed of sub-capacities to represent London, Paris, New York, and the relation of being north of, along with general capacities for predication and conjunction. All these sub-capacities must be exercised and appropriately coordinated in order for the Mentalese sentence to be instantiated.

A complete account should develop my intuitive formulations more systematically, including provision of a compositional semantics. A complete account should also address additional logical compounding devices, including disjunction, negation, the conditional, and the quantifiers. There are many delicate philosophical and technical details here.[6] My goal is not to provide a complete account but rather to indicate a promising direction for future research.

I have focused so far on representational complexity as it arises within high-level cognition. Some authors propose that high-level cognition has a different format than *perception*

[6] For example, one must extend (Δ) with a generalized clause stipulating what it is for mental representation *R* to have semantic value α.

(Fodor, 2008; Peacocke, 1992), *mental imagery* (Kosslyn, 1980), *analogue magnitude representation* (Beck, 2012), or other relatively low-level representational phenomena. For example, Burge (2010b) holds that perceptual representations exhibit nothing like the logical structure characteristic of high-level thought. In (Rescorla, 2009), I suggested that something similar may hold of cognitive maps. These issues remain murky and controversial. I think that C-RTM can help. To elucidate representational format, we should analyze how complex representational capacities arise from the appropriate joint exercise of representational sub-capacities. Different representational formats correspond to differently structured ways of exercising sub-capacities. I acknowledge that my formulations are sketchy. Once again, I offer them only to indicate a promising path forward.

Even in its present preliminary state, C-RTM validates the traditional thought that mental representations can combine to form complex mental representations. There is a natural sense in which type 𝕊𝕊𝕊𝕊𝟘 results from appropriately combining types 𝟘 and 𝕊: 𝕊𝕊𝕊𝕊𝟘 is instantiated only when 𝟘 and 𝕊 are instantiated in an appropriate way. Similarly, there is a natural sense in which 𝕃𝕠𝕟𝕕𝕠𝕟 𝕚𝕤 𝕟𝕠𝕣𝕥𝕙 𝕠𝕗 ℙ𝕒𝕣𝕚𝕤 results from appropriately combining 𝕃𝕠𝕟𝕕𝕠𝕟, ℙ𝕒𝕣𝕚𝕤, and 𝕚𝕤 𝕟𝕠𝕣𝕥𝕙 𝕠𝕗. Just as a C-major chord results when a musician appropriately combines individual notes, and just as the Biles results when a gymnast appropriately combines individual gymnastics moves, so does a complex mental representation result when the mind appropriately combines mental representations. Complex types are rightly so-called because they mark the structured instantiation of other types. Complex mental representations are rightly so-called because they mark the structured instantiation of other mental representations.

§5.5 Internally structured mental representations?

Philosophers usually construe the complexity of mental representations more literally than I have done. Most discussants hold that a complex mental representation is a structured entity composed of less complex mental representations. Call this *the mereological thesis*. According to the mereological thesis, a complex mental representation has other mental representations as literal parts. I remain agnostic regarding the mereological thesis, as befits my agnosticism regarding whether complex types have internal structure. What matters most is that a complex type marks the structured instantiation of incorporated types. Specifically, a complex mental representation marks the structured exercise of representational capacities. Whether the mental representation itself has structure is not nearly so important.[7]

My agnosticism regarding the mereological thesis mandates a slight departure from the usual treatment of compositionality. Philosophers usually gloss compositionality in mereological terms. They say, roughly, that the meaning of a complex mental representation is determined by the meanings of its *parts* and the way in which those parts are put together. A mereological gloss is unavailable to me, since I do not assume that mental representations have parts. Instead, I say that the meaning of a complex representation is determined by the meanings of the representations that it incorporates and the way that it incorporates them. As illustrated by my discussion of the toy mathematical language, we need not assume that a complex mental representation has internal structure when specifying its compositional semantics. We can instead assume that it marks the structured instantiation of other mental representations.

[7] Burge (2009, 2010a) posits an array of mental representations employed in thought, perception, navigation, and other mental activities. He identifies these items with *mental representational contents* (2009, p. 248). In many respects, my position is similar to Burge's. One difference is that Burge regards mental representations as structured. He writes: "At bottom, representational contents are just kinds, or aspects of kinds, of psychological states. The structure of representational contents marks structural aspects of the capacities embodied in the psychological states" (2010a, p. 41). In contrast, I remain neutral as to whether mental representations have internal structure. This difference is not as significant as it may initially appear, because I still recognize an important sense in which complex mental representations result from combining together other mental representations.

Proponents offer various arguments for the mereological thesis (Davis, 2003, pp. 368-406). Most famously, Fodor and Pylyshyn (1988) highlight a phenomenon called *systematicity*: there are systematic relations among which thoughts a thinker can entertain. To use their example, someone who can think *that John loves the girl* can also think *that the girl loves John*. Fodor and Pylyshyn write (1988, p. 39):

> The systematicity of thought shows that there must be structural relations between the mental representation that corresponds to the thought that John loves the girl and the mental representation that corresponds to the thought that the girl loves John; namely, the two mental representations…*must be made of the same parts*. But if this explanation is right (and there don't seem to be any others on offer), then mental representations have internal structure.

Thus, Fodor and Pylyshyn defend the mereological thesis by arguing that it underwrites the best explanation for systematicity.

Let us consider in more detail how the explanation is supposed to go. Assume that Fred can think *that John loves Mary*. Why can he also think *that Mary loves John*? Presumably Fodor and Pylyshyn have in mind an explanation along the following lines:

(7) If Fred can think *that John loves Mary*, then he can stand in an appropriate relation *T* to a structured mental representation *R*. The parts of *R* can be recombined to form a different structured mental representation, corresponding to the thought *that Mary loves John*. Fred can also stand in relation *T* to that second structured mental representation, since he has access to all its parts. Thus, he can think *that Mary loves John*.

I agree that this is one possible explanation. But an alternative explanation is possible:

(8) When Fred thinks *that John loves Mary*, he exercises a complex representational capacity *C* composed of sub-capacities that include: a capacity for predication; a capacity to represent John; a capacity to represent the loving relation; and a capacity to represent Mary. Because Fred has all these sub-capacities, he also has a second complex representational capacity *C**, corresponding to the thought *that Mary loves John*. He has capacity *C** because *C** involves precisely the same sub-capacities as *C*. The sub-capacities are merely coordinated in a different way when one exercises *C* versus *C**. Thus, Fred can think *that Mary loves John*.

I think that (8) seems as good an explanation as (7).Yet (8) does not assume that complex mental representations are structured. It does not even mention mental representations. A nominalist who refuses to reify mental representation types could accept (8). Of course, I am not a nominalist. I cheerfully reify, correlating each representational capacity mentioned by (8) with a mental representation. But I do not claim that reification improves the explanation given by (8). Structural relations among representational capacities, not structural relations among mental event types, explain why Fred's ability to think *that John loves Mary* entails his ability to think *that Mary loves John*. Contrary to Fodor and Pylyshyn, systematicity does not show that there are structural relations among mental representations. It only shows that there are structural relations among representational capacities. Thus, systematicity provides no support for the mereological thesis over my less committal viewpoint.[8]

Fodor and Pylyshyn ask (1988, p. 44): "*how could* the mind be so arranged that the ability to be in one representational state is connected with the ability to be in others that are

[8] Here I build on Evans's discussion. Evans (1982, p. 104) endorses a form of systematicity that he calls the *Generality Constraint*. To explain why thinkers satisfy the Generality Constraint, Evans cites the structured nature of representational capacities. He also writes (1982, p. 101): "I should prefer to explain the sense in which thoughts are structured, not in terms of their being composed of several distinct *elements*, but in terms of their being a complex of the exercise of several distinct conceptual *abilities*."

semantically nearby? What account of mental representation would have this consequence?". I reply that the ability to be in a representational state is typically a complex capacity composed of sub-capacities; the sub-capacities can be redeployed so that the thinker instantiates "semantically nearby" states. The crucial observation is that mental representation involves the structured exercise of redeployable representational capacities. This observation concerns capacities, not the metaphysics of types. *Even if we were to conclude that complex mental representations have mereological structure*, it would still be preferable to explain systematicity without invoking that structure. The best explanation would still adduce the structured way that mental events deploy representational capacities, without any detour through the structure of mental event types. The former kind of structure, not the latter, is explanatory fundamental.

I have critiqued Fodor and Pylyshyn's systematicity argument for the mereological thesis. Parallel considerations apply to other well-known arguments for the thesis. Parallel considerations show that any mental phenomena customarily explained by the mereological thesis are explained as well, if not better, by citing structural relations among representational capacities. For present purposes, I must leave my assessment undefended.

§6. Mode of presentation

In the previous section, I invoked representational capacities to clarify the complexity of mental representations. I now invoke them to clarify another phenomenon widely recognized among proponents of RTM: distinct mental representations may share the same denotation.

Frege (1892/1997) argued that a thinker can represent a single denotation in different *ways*, or under different *modes of presentation*. He illustrated by considering a thinker who believes that Hesperus is Hesperus but does not believe that Hesperus is Phosphorus. Frege says that the thinker mentally represents the same denotation (Venus) under two distinct modes of presentation. But what exactly are modes of presentation? Here are two well-known proposals:

- Fodor (1994) proposes that we gloss modes of presentation as mental representations. For example, we can posit distinct, co-referring Mentalese words that denote Venus.
- Evans (1982, pp. 100-105), following Geach (1957), proposes that we elucidate modes of presentation in terms of *abilities*. He writes: "[w]hen two thought-episodes depend on the *same* ability to think of something, we can say that the thing is thought about *in the same way*" (p. 101). Conversely, the thing is thought about in *different* ways when the two thought-episodes depend on different abilities to think about it.[9]

C-RTM allow us to combine Fodor's proposal with Evans's. We can say that distinct modes of presentation are distinct mental representations, marking the exercise of distinct representational capacities. For example, we may posit a Mentalese word **Hesperus** that marks the exercise of a capacity to represent Venus in higher-level thought and a distinct Mentalese word **Phosphorus** that marks the exercise of a *different* capacity to represent Venus. Distinct co-referring mental representations correspond to distinct capacities for representing a single denotation.

Why distinguish the Mentalese words **Hesperus** and **Phosphorus**? Why say that the corresponding representational capacities differ? To simplify matters, let us consider a thinker who has never examined the heavens closely enough to perceive Venus directly. The thinker is still able to represent Venus within her thought. She acquires a capacity to represent Venus when she learns the English word "Hesperus." This capacity emerges as she masters those aspects of linguistic practice involving the word "Hesperus." When she learns the English word "Phosphorus," she acquires a second capacity to represent Venus. This second capacity emerges

[9] Citing Evans as an influence, Beck (2013) likewise elucidates modes of presentation in terms of abilities.

as she masters those aspects of linguistic practice involving the word "Phosphorus." That the two capacities are different is evidenced by several facts:

- They are acquired on different occasions and through exposure to different aspects of linguistic practice.
- A thinker can acquire one without acquiring the other.
- Even after a thinker has acquired both, she can exercise one without exercising the other (e.g. she can think of an entity *as Hesperus* without thinking of it *as Phosphorus*).

Hesperus and **Phosphorus** are distinct types that mark the exercise of distinct capacities to represent the same planet. The thinker instantiates the first type when she represents Venus via her connection to our linguistic practices surrounding the word "Hesperus." She instantiates the second type when she represents Venus via her connection to our linguistic practices surrounding the word "Phosphorus."[10]

Modes of presentation are crucial for understanding mental representation in general, not just high-level thought. A good illustration is *sensory cue combination*. The perceptual system typically estimates distal conditions based upon multiple cues. It can estimate an object's size based on visual feedback or haptic feedback. It can estimate depth using binocular disparity, motion parallax, convergence, and many other visual cues. Bayesian perceptual psychology offers detailed models of cue combination (Trommershäuser, Körding, and Landy, 2011). A key presupposition underlying some of the most successful models is that distinct sensory cues correspond to distinct co-referring perceptual representations (Rescorla, forthcoming b). For

[10] If the thinker directly observes Venus, then she may on that basis acquire a third capacity to represent Venus within thought. This capacity is distinct from the capacities corresponding to **Hesperus** and **Phosphorus**, because it is grounded in perception of Venus rather than linguistic competence. We should therefore postulate a third mental representation that marks the exercise of this third representational capacity.

example, the perceptual system employs a vision-based representation that denotes distal size *s*, and it employs a distinct touch-based representation that also denotes *s*. Distinct co-referring perceptual representations denote the same denotation, but they do so in different *ways*.

Proponents of C-RTM can say that distinct co-referring perceptual representations mark the exercise of distinct capacities to represent the same denotation. For example, a vision-based representation that denotes size *s* marks the exercise of a different capacity than a touch-based representation that denotes *s*. The capacities are tied to different sensory cues. A perceiver has the first capacity only if she has (or is appropriately related to) a perceptual system wired to estimate distal size based upon visual cues. A perceiver has the second capacity only if she has (or is appropriately related to) a perceptual system wired to estimate distal size based upon haptic cues. A perceptual system might be wired the first way without being wired the second way, and vice-versa. Hence, the capacities are distinct. Similarly for other cases where distinct perceptual representations correspond to distinct sensory cues.

A mental event represents denotation *d* only because the event occurs within the mental activity of a creature with representational capacities. The creature is able to represent *d* by virtue of past causal interactions with *d*, or evolutionary relations to progenitors that interacted with *d*, or embedding within a linguistic practice that represents *d*, or internal psychological processing, or some combination of these and possibly other factors. *Some* combination of factors renders the creature able to represent *d*. A relevantly different combination of factors yields a different capacity to represent *d*. We register such differences by positing distinct co-referring mental representations. Thus, C-RTM grounds mental co-reference in a fundamental feature of mental representation: diverse factors may render a creature able to represent a single denotation.

One might question how I am individuating representational capacities. Why individuate them in so fine-grained a way? There is such a thing as the sheer capacity to represent Venus. You exercise that capacity when you think *that Hesperus has craters* and also when you think *that Phosphorus has craters*. Why not emphasize the general representational capacity deployed by both thoughts? Likewise, why not emphasize the general representational capacity to represent size *s* within perception, without distinguishing between vision-based and touch-based exercises of the capacity? Why not adopt a coarse-grained scheme that individuates representational capacities entirely through the represented denotations?

I reply that we should taxonomize mental events so as to support good explanations. Hesperus-thoughts and Phosphorus-thoughts play different roles in cognition. They figure differently within belief-fixation (certain astronomical observations may lead you to believe *that Hesperus has craters* but not *that Phosphorus has craters*), decision-making (you may make different plans if you want to test *whether Hesperus has craters* than if you want to test *whether Phosphorus has craters*), linguistic comprehension (subpersonal linguistic processing proceeds differently if someone tells you *that Hesperus has craters* than if someone tells you *that Phosphorus has craters*), and other mental processes. The differences diminish markedly if you learn *that Hesperus is Phosphorus*, but even then some differences persist, especially differences in linguistic processing. A satisfactory theory must track the differences. To do so, it must adopt a fine-grained taxonomic scheme that differentiates Hesperus-thoughts from Phosphorus-thoughts. A coarse-grained taxonomic scheme that recognizes only the sheer capacity to represent Hesperus will not serve nearly as well as a finer-grained scheme that distinguishes among capacities for representing Hesperus. The coarse-grained scheme may be useful for some purposes. Overall, it is less apt to promote fruitful explanation.

How exactly do distinct, co-referring mental representations differ? For example, under what circumstances does a mental event instantiate the type **Hesperus** rather than the type **Phosphorus**? The literature on RTM has addressed these questions extensively (Field, 2001, pp. 55-58; Fodor, 1994; Schneider, 2011; Stich, 1983). Lying in the background is a widespread conviction that we should admit entities into our discourse only when we can associate them with well-defined *identity conditions*, i.e. conditions for re-identifying an entity as the same again. Quine (1969, p. 23) forcefully argues as much, summarizing his viewpoint through the slogan "no entity without identity." As applied to mental representations, Quine's viewpoint requires that we specify conditions under which mental representation *R* is the same as mental representation *S*. Adopting this viewpoint, proponents of RTM seek a plausible individuative scheme for mental representations, while opponents often contend that existing individuative schemes are unsatisfactory and hence that RTM is problematic (Prinz, 2011).

C-RTM holds that different mental representations correlate with different representational capacities. However, this just pushes the bump under the rug from individuation of mental representations to individuation of representational capacities. Under what conditions are mental events *e* and *e** exercises of the same representational capacity? For example, what exactly distinguishes the representational capacity correlated with **Hesperus** from the representational capacity correlated with **Phosphorus**? Lacking answers to these questions, we still lack non-circular identity conditions for mental representations.

My response is to deny that we must provide non-circular identity conditions for entities before admitting them into our discourse. Quotidian and scientific discourse posit diverse entities: properties, events, persons, words, species, symphonies, cities, etc. Outside the realm of extensional mathematics taken by Quine as a paradigm, we can rarely supply anything like non-circular identity conditions. Identity conditions are helpful when available, but they are not mandatory. This is true for entities in general, and it is true for types specifically. Taxonomic schemes found in quotidian and scientific discourse rarely come associated with explicit necessary and sufficient conditions. To pick only one example, we have nothing like non-circular identity conditions for natural language words (e.g. Bromberger, 2011; Hawthorne and Lepore, 2011). Few philosophers would argue on that basis for banishing natural language words from our ontology. Likewise, non-circular identity conditions are not needed for mental representations to play a valuable role within psychological theorizing. Cognitive science practice demonstrates that we can make substantial explanatory progress by reifying mental representations absent non-circular identity conditions.

Luckily, one can illuminate how mental representations are individuated without furnishing non-circular identity conditions. As a fairly straightforward example, consider the co-referring representations $\mathbb{SSSS0}$ and $\mathbb{SS0}+\mathbb{SS0}$. These types mark the respective exercise of complex representational capacities. We can say pretty explicitly how the capacities differ, because we can describe how each capacity decomposes into the exercise of simpler capacities. More generally: once we identify a mental representation as complex, we can usually then analyze the complex representational capacity correlated with it.

Whether or not a mental representation is complex, we can often say something helpful about the corresponding representational capacity. A few examples:

- *Singular terms*. Over the past century, we have learned a lot about representational capacities corresponding to Mentalese singular terms. Most importantly, Kripke (1980) shows that you can have such a capacity without knowing descriptive information that distinguishes the denotation from other possible denotations. You can think about Kurt Gödel *as Kurt Gödel* without entertaining or grasping any definite description that discriminates Kurt Gödel from other people.

- *Predicates*. You can think about elm trees *as elm trees* even though you have no knowledge that differentiates elm trees from beech trees (Putnam, 1975b). You can think about arthritis *as arthritis* even if you think arthritis is a disease that occurs in muscles rather than joints (Burge, 2007). In general, you can exercise the representational capacity corresponding to a Mentalese predicate even if you lack discriminating information about the predicate's extension, and despite significant false beliefs about the extension, so long as you are suitably embedded in an appropriate linguistic practice.
- *Perceptual representations*. Perceptual psychology sheds considerable light upon the representational capacities deploying during perception. Consider a vision-based perceptual representation of size *s*. As I argue elsewhere (Rescorla, forthcoming b), you can instantiate the representation in response to a wide range of possible retinal stimulations and despite large changes in Bayesian priors. The corresponding representational capacity is tied to your perceptual system's general capacity for estimating size based on one or more visual cues, not to the specific retinal stimulations that serve as inputs to visual estimation or to the specific Bayesian priors deployed during perceptual processing.

Obviously, these observations fall far short of non-circular identity conditions. Nevertheless, as such observations accrue, we gradually clarify how mental representations and their affiliated representational capacities should be individuated.

In my opinion, individuation of representational capacities is best elucidated on a case by case basis, by interrogating folk psychology or cognitive science regarding specific capacities. Insight is more likely to emerge from detailed study of particular examples than from grand attempts at overarching non-circular identity conditions. The present paper aims not to supply satisfactory elucidations but rather to delineate a framework within which elucidation can occur.[11]

§7. An individuative role for representation

§§5-6 argued that C-RTM sheds light upon some commitments that are widespread among proponents of RTM. I now show that C-RTM diverges in at least one important respect from other contemporary versions of RTM.

Say that an entity is *semantically indeterminate* when it does not have its meaning essentially. A semantically indeterminate entity could have had a different meaning without any change in its fundamental nature, identity, or essence. Say that an entity is *semantically neutral* when it bears an arbitrary relation to its meaning (assuming it even has meaning). A semantically neutral entity could have had *arbitrarily different* meaning, or no meaning at all, without any change in its fundamental nature, identity, or essence. Semantic neutrality entails semantic indeterminacy, but not vice-versa: semantic indeterminacy entails that the entity could have had *some* different meaning, while semantic neutrality entails that it could have had *any* different

[11] Because my approach invokes modes of presentation, it accommodates many phenomena that modes of presentation are famously well-suited to handle. Consider *reference failure*. Suppose we want to describe Le Verrier's mental state when he conjectured *that Vulcan orbits between Mercury and the sun*. (In fact, there is no such planet as Vulcan.) We may posit a Mentalese word **Vulcan** that lacks any denotation. This Mentalese word marks the exercise of a defective representational capacity. The capacity is *representational* because it is a capacity to attempt reference to an object. The capacity is *defective* because the attempt fails: someone who exercises the capacity does not thereby succeed in referring to any object. Developing these remarks and bringing them into contact with the large literature on reference failure are tasks for another paper.

meaning. Most communal representations are semantically neutral. For example, the word "dog" means dog, but it could just as well have meant cat, or anything else, or nothing at all.

Over the past few decades, many philosophers have pursued a semantically indeterminate taxonomic scheme for mental representations. This approach originates with Fodor (1981). He holds that Mentalese expressions have *formal syntactic* properties, and he introduces an array of corresponding formal syntactic Mentalese types. A Mentalese syntactic type has representational import, but it does not have its representational import *essentially*. For example, Fodor posits a Mentalese word DOG that denotes dogs. According to Fodor, DOG could have had a different denotation had it played a different role in the thinker's psychological activity. In his early writings, Fodor regarded formal syntax as semantically indeterminate but not semantically neutral. He claimed that Mentalese syntactic type *constrains* meaning while leaving meaning underdetermined (1981, pp. 225-253). Fodor's later writings (1994, 2008) suggest the stronger thesis that Mentalese syntactic types are semantically neutral. Many authors explicitly advocate a semantically neutral taxonomic scheme for Mentalese syntax (e.g. Egan, 1992, p. 446; Field, 2001, p. 58; Haugeland, 1985, p. 91, pp. 117-123; Pylyshyn, 1984, p. 50).

The contemporary consensus in favor of semantic indeterminacy departs from philosophical tradition. Historically, proponents of RTM tended to individuate mental representations in representational terms. For example, Ockham does not hold that one can "hive off" a mental word's semantic or representational properties and leave behind a theoretically significant formal syntactic residue. Philosophers did not traditionally regard mental representations as items subject to reinterpretation in the way that communal representations are subject to reinterpretation.

Say that an entity is *semantically permeated* when it is not semantically indeterminate. A semantically permeated entity is not a piece of formal syntax requiring an interpretation. Rather, its semantics is built into its inherent nature. It "comes with its meaning attached." I will argue that mental representations as construed by C-RTM are semantically permeated.

§7.1 Semantically permeated mental types

According to C-RTM, a mental representation is a type that marks the exercise of a representational capacity. The type is instantiated only when the thinker (or a mental subsystem) exercises the capacity. Tokens of the type must have whatever representational properties are implicated by that exercise. Thus, mental representations are individuated at least partly through their representational properties. Examples:

- *Mentalese words*. The Mentalese word **dog** denotes dogs: it marks the exercise of a capacity to represent dogs within high-level thought. A mental event instantiates **dog** only if it is an exercise of that capacity, to be which the event must represent dogs. Thus, **dog** is not an uninterpreted formal item that could just as easily have denoted cats. A mental event that instantiates **dog** must be an exercise of the corresponding representational capacity, so it must represent dogs rather than cats. The Mentalese word, by its inherent nature, already mandates one specific interpretation.
- *Perceptual representations*. Consider a perceptual representation *S* that marks the exercise of a capacity to represent sphericality. A perceptual state instantiates *S* only if it is an exercise of the corresponding capacity, to be which it must represent sphericality. So *S* could not have denoted another shape, let alone some other distal

property. *S*, by its inherent nature, is instantiated by only by mental states that represent sphericality.

- *Cognitive maps*. A cognitive map *M* marks the exercise of a capacity to represent some spatial layout. In order for *M* to be instantiated, the animal must exercise the correlated capacity. *M* could not have represented a different spatial layout, nor could it have had non-spatial representational import. *M*, by its inherent nature, is instantiated only by mental states that represent a particular spatial layout.

Similarly, the mental numerals from §5.3 have built-in denotations, as specified by the compositional semantics. In general, mental representations are not subject to arbitrary reinterpretation. They are imbued with representational import by their inherent natures. My semantically permeated approach accords well with much of the historical literature even while it flouts current conventional wisdom.[12]

Semantically permeated types are individuated through their representational properties, but not *all* representational properties need play an individuative role. For example, a Mentalese sentence that has its truth-*condition* essentially need not have its truth-*value* essentially. **London is north of Paris** is true iff London is north of Paris, and this truth-condition is inherent to the Mentalese sentence. Whether the Mentalese sentence is true, on the other hand, depends on how the world is. It depends on whether London is north of Paris, which is not essential to the type. Whether London is north of Paris plays no role in individuating the complex type.

For a subtler example where representational properties need not play an individuative role, consider *indexicality*. Looking at a cube, I may judge *that that cube is blue* and intend *to grab that cube*. My judgment's truth-condition depends on the specific contextually-determined

---

[12] Burge pursues a similar approach, applied especially to concepts (2007, p. 292) and perceptual representations (2010a, p.76).

cube, as does my intention's fulfillment-condition. Indexicality also arises pervasively in perception (Burge, 2010a). If my perceptual state represents a perceived cube as blue, then the state's accuracy-condition depends upon the specific contextually-determined cube.

To focus the discussion, let us consider two thinkers *A* and *B* who are psychological duplicates except that they perceive distinct, qualitatively indiscernible cubes C_A and C_B. *A* intends to grab C_A, while *B* intends to grab C_B. The two intentions are the same in all relevant respects except that they represent distinct contextually-determined cubes. Suppose we posit a *mental demonstrative* that figures in *A*'s intention. How should we individuate this demonstrative? Evans (1982) and McDowell (1998) individuate mental demonstratives in an *object-dependent* way, so that *A*'s mental demonstrative is different from *B*'s corresponding mental demonstrative. Burge (2005) argues that we should sometimes individuate mental demonstratives in an *object-independent* way, so that *A* and *B* instantiate the same mental demonstrative type.

The main point I want to stress is that C-RTM can accommodate both the object-dependent and the object-independent viewpoints. One can individuate representational capacities in either an object-dependent or object-independent fashion. From an object-dependent viewpoint, *A* has a capacity to represent cube C_A, while *B* has a different capacity to represent a different cube C_B. From an object-independent viewpoint, *A* and *B* share a common capacity to represent some perceptually presented cube, and this common capacity determines different denotations when exercised in different contexts. If we adopt the object-dependent viewpoint, we posit a mental demonstrative that comes with its specific denotation attached. If we adopt the object-independent viewpoint, we posit a mental demonstrative that comes with certain representational properties attached (e.g. the property of being a mental demonstrative), not including its contextually-determined denotation. C-RTM allows *both* viewpoints. It also allows theorists to adopt the first viewpoint for some projects and the second viewpoint for other projects. Thus, C-RTM is equally friendly to the object-dependent and object-independent individuative schemes. Semantic permeation does not entail that a mental representation has its context-dependent representational properties essentially.[13]

§7.2 Representation as explanatorily central

Why do I adopt a semantically permeated taxonomic scheme for mental representations? Why individuate mental representations in representational terms? The main reason is that I want to track how explanation proceeds within cognitive science.

Mental representations are abstract entities whose primary role in our discourse is to facilitate taxonomization of mental events for explanatory purposes. When we decide how to individuate these entities, explanation should be our main touchstone. How we taxonomize mental events affects which explanations we can provide, so we should choose a taxonomic scheme that underwrites good explanations. Moreover, our best guide to good psychological explanation is actual explanatory practice within scientific psychology. As I urged in §2, numerous areas of cognitive science classify mental events through their representational properties. Representation occupies explanatory center stage in current scientific theories of perception, motor control, mammalian navigation, high-level cognition, linguistic communication, and numerous other core mental phenomena. The theories cite representational properties so as to sort mental events into types. I reify the types by positing semantically permeated mental representations. I thereby codify current cognitive science practice in

[13] An object-independent version of C-RTM requires that we revise (Δ) to allow for context-dependent denotations. One option is to emend (Δ) along the following lines: R denotes d in context γ iff R marks the exercise of some representational capacity C and any exercise of C in context γ is a mental event that represents d.

ontologically loaded terms. In some areas, such as Bayesian perceptual psychology, the science already posits semantically permeated mental representations.

Of course, one might employ a semantically neutral taxonomic scheme *in addition to* a semantically permeated taxonomic scheme. One might classify mental events in representational terms for certain explanatory purposes but not for other explanatory purposes. To what extent does cognitive science practice actually involve non-representational taxonomization of representational mental events?

Non-representational description figures crucially in *neuroscience*. Neuroscientists frequently adduce firing rates, action potentials, and other such non-representational aspects of neural activity. A typical neural event type is semantically neutral. It could have had any arbitrarily different representational import (or none at all) depending on its role in the cognitive system as a whole. Neurophysiological description leaves open which if any representational properties neural events have. However, philosophers who espouse semantically indeterminate mental representations do not usually envisage a neurophysiological taxonomic scheme. Building on Putnam's (1975b) critique of type-physicalism, these philosophers think that psychological description should be *multiply realizable* in the neural. They think that psychological description should admit wildly different neural instantiations. Mentalese syntax, like psychological description more generally, is supposed to be multiply realizable. Accordingly, advocates of formal mental syntax pursue a taxonomic scheme for mental representations that does not cite neural properties but instead cites multiply realizable psychological properties (Fodor, 2008, p. 91; Haugeland, 1985, p. 5; Stich, 1983, p. 151).

In my opinion, current cognitive science does not support any such formal syntactic taxonomic scheme (Rescorla, 2017b). The proposed scheme plays no role within current scientific theories of perception, motor control, mammalian navigation, or numerous other core mental processes. Researchers describe these processes in representational terms. They also try to illuminate how representational mental activity is grounded in underlying neural activity. Researchers describe the processes through *multiply realizable representational descriptions* and *non-representational descriptions that are not multiply realizable*. They do not employ *multiply realizable non-representational descriptions*. For example, perceptual psychology describes perceptual inference in representational terms *as opposed to* formal syntactic terms. (Cf. Burge, 2010a, pp. 95-97.) Based on current scientific practice, I see little evidence that we can "hive off" a mental representation's representational import and isolate an explanatorily significant formal syntactic residue. Perhaps cognitive science as it evolves will eventually individuate mental representations in formal syntactic fashion. Current science provides little reason to expect so.[14]

I am skeptical about formal syntactic description of *representational* mental events, not formal syntactic description of mental events more generally. Cognitive scientists deploy multiply realizable non-representational description to explain some mental phenomena, such as certain kinds of low-level insect navigation (Rescorla, 2013). When regimenting these explanations, it is natural to postulate semantically indeterminate mental event types. However,

[14] An exception is Carey's (2009) work on concept acquisition, which postulates something like formal syntactic Mentalese types. On Carey's approach, a child can acquire a concept through exposure to a new "explicit symbol," such as a natural language word. "The capacity for explicit symbolization makes possible the creation of mental symbols that are not yet connected to anything in the world… [M]ental symbols are established that correspond to newly coined or newly learned explicit symbols. These are initially placeholders, getting whatever meaning they have from their interrelations with other symbols" (Carey, 2009, p. 474). Through analogical reasoning, induction, and other techniques, the child "bootstraps" her way from placeholder symbols to new concepts. Placeholder symbols are individuated in semantically indeterminate (perhaps even semantically neutral) fashion. They are uninterpreted items that become endowed with denotations only during the bootstrapping process. If Carey is right, then formal syntactic taxonomization should play a central role in any complete theory of concept acquisition. However, Carey's approach is controversial (e.g. Rips and Hespos, 2011), and the crucial notion of "bootstrapping" remains murky. I doubt that scientific research into concept acquisition (as opposed to perception, motor control, navigation, linguistic comprehension, causal learning, and various other psychological domains) is sufficiently developed to support solid conclusions about the individuation of mental representations.

the relevant mental events are not representational. They do not represent the world as being a certain way. I doubt that formal syntactic description of *representational* mental events offers any explanatory benefit to cognitive science explanation.

The central issue here is *explanation*, not *existence*. No doubt we can describe representational mental events in the formal syntactic terms favored by Fodor. We can also posit semantically indeterminate types corresponding to a formal syntactic taxonomic scheme. What I question is whether the scientific study of mental representation gains any explanatory value by positing such types or by employing the taxonomic scheme that they embody.

Over the past few decades, philosophers have advanced various arguments that any complete cognitive science requires multiply realizable, semantically indeterminate descriptions of representational mental events. I believe that all these arguments fail. I have critiqued the most prominent arguments elsewhere (Rescorla, 2017a).

Philosophers commonly emphasize a sharp distinction between representational *vehicles* and representational *contents*. For example, Fodor (1981, 1994, 2008) regards Mentalese syntactic types as vehicles that express contents. It may seem that the items I am calling "mental representations" should not be so-called, because they are closer to contents than to vehicles. In particular, their semantically permeated character may seem to render them more "content-like" than "vehicle-like." Personally, I see no great import in whether we classify mental representations construed in my terms as vehicles, contents, both, or neither. Excessive focus upon the locutions "vehicle" and "content" has deflected attention from the more fundamental question of which taxonomic schemes best serve psychological explanation. Mental representations as I construe them are perfectly tailored to serve good psychological explanation, because they are directly grounded in the representational capacities already presupposed by our best cognitive science explanations.

§8. A framework for studying mental representations

C-RTM vindicates mental representations as scientifically respectable and metaphysically harmless. It clarifies how they combine to form complex representations. It elucidates differences among co-referring mental representations. Most importantly, it promotes a traditional semantically permeated perspective on their individuation. By linking mental representations to the exercise of representational capacities, it honors the pivotal role that representational properties play within psychological explanation. C-RTM assigns representationality its rightful centrality within the philosophical study of mental representations, dispelling the explanatorily idle formal syntactic properties that typically mar contemporary expositions.

I have offered a framework for inquiry, not a finished theory. The framework raises numerous issues that merit further scrutiny. In future work, I will enhance the framework with supplementary theoretical reflections and more detailed case studies.

Acknowledgments

I am grateful to Jacob Beck, Eric Mandelbaum, Marcus Giaquinto, Mark Sainsbury, and two anonymous referees for this volume for comments that greatly improved the paper. I presented versions of this material at the University of Pittsburgh Center for Philosophy of Science; the Institute of Philosophy, London; Ohio State University; the University of California, Irvine; and the 2018 Workshop on Perceptual Capacities and Psychophysics at Rutgers

University. I thank the audience members on those occasions for their helpful feedback, especially Casey O'Callaghan, David Papineau, Christopher Peacocke, Jake Quilty-Dunn, Richard Samuels, Nicholas Shea, and Declan Smithies. My research was supported by a fellowship from the National Endowment for the Humanities. Any views, findings, conclusions, or recommendations expressed in this publication do not necessarily reflect those of the National Endowment for the Humanities.

Works Cited

Armstrong, D. 1980. *A Theory of Universals: Universals and Scientific Realism, vol. 2*. Cambridge: Cambridge University Press.

Beck, J. 2012. "The Generality Constraint and the Structure of Thought." *Mind* 121: pp. 563-600.

---. 2013. "Sense, Mentalese, and Ontology." *Protosociology* 30: pp. 29-48.

Bennett, K. 2013. "Having a Part Twice Over." *Australasian Journal of Philosophy* 91: pp. 83-103.

Bromberger, S. 2011. "What Are Words? Comments on Kaplan (1990), on Hawthorne and Lepore, and on this Issue." *Journal of Philosophy* 108: pp. 486-503.

Burge, T. 2005. "Disjunctivism and Perceptual Psychology." *Philosophical Topics* 33: pp. 1-78.

---. 2007. *Foundations of Mind*. Oxford: Oxford University Press.

---. 2009. "Five Theses on *De Re* States and Attitudes." In *The Philosophy of David Kaplan*, eds. J. Almog and P. Leonardi. Oxford: Oxford University Press.

---. 2010a. *Origins of Objectivity*. Oxford: Oxford University Press.

---. 2010b. "Steps toward Origins of Propositional Thought." *Disputatio* 4: pp. 39-67.

Carey, S. 2009. *The Origin of Concepts*. Oxford: Oxford University Press.

Chemero, A. 2009. *Radical Embodied Cognitive Science*. Cambridge: MIT Press.

Churchland, P. M. 1981. "Eliminative Materialism and the Propositional Attitudes." *Journal of Philosophy* 78: pp. 67-90.

Davis, W. 2003. *Meaning, Expression, and Thought*. Cambridge: Cambridge University Press.

---. 2014. "On Occurrences of Types in Types." *Australasian Journal of Philosophy* 92: pp. 349-363.

Egan, F. 1992. "Individualism, Computation, and Perceptual Content." *Mind* 101: pp. 443-459.

Evans, G. 1982. *The Varieties of Reference*. Oxford: Clarendon Press.

Field, H. 2001. *Truth and the Absence of Fact*. Oxford: Clarendon Press.

Fodor, J. 1975. *The Language of Thought*. New York: Thomas Y. Crowell.

---. 1981. *Representations*. Cambridge: MIT Press.

---. 1987. *Psychosemantics*. Cambridge: MIT Press.

---. 1994. *The Elm and the Expert*. Cambridge: MIT Press.

---. 2008. *LOT2*. Oxford: Clarendon Press.
Fodor, J., and Pylyshyn, Z. 1988. "Connectionism and Cognitive Architecture: A Critical Analysis." *Cognition* 28: pp. 3-71.
Frege, G. 1892/1997. On *Sinn* and *Bedeutung*. Rpt. *The Frege Reader*, ed. M. Beaney, trans. M. Black. Malden: Blackwell.
Gallistel, C. R. 1990. *The Organization of Learning*. Cambridge: MIT Press.
---. 1999. "Coordinate Transformations in the Genesis of Action." In *Cognitive Science: Handbook of Perception and Cognition*, 2nd ed., eds. B. Bly and D. Rumelhart. New York: Academic Press.
Gallistel, C. R. and Gelman, R. 2005. "Mathematical Cognition." In *The Cambridge Handbook of Thinking and Reasoning*, eds. K Holyoak and R. Morrison. Cambridge: Cambridge University Press.
Geach, P. 1957. *Mental Acts*. London: Routledge and Paul.
Haugeland, J. 1985. *Artificial Intelligence: The Very Idea*. Cambridge: MIT Press.
Hawley, K. 2010. "Mereology, Modality and Magic." *Australasian Journal of Philosophy* 88: pp. 117–33.
Hawthorne, J., and Lepore, E. 2011. "On Words." *Journal of Philosophy* 108: pp. 447-485.
Heim, I., and Kratzer, A. 1998. *Semantics in Generative Grammar*. Malden: Blackwell.
Helmholtz, H. von. 1867. *Handbuch der Physiologischen Optik*. Leipzig: Voss.
Knill, D., and Richards, W., eds. 1996. *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.
Kripke, K. 1980. *Naming and Necessity*. Cambridge: Harvard University Press.
Kosslyn, S. 1980. *Image and Mind*, Cambridge: Harvard University Press.
Lewis, D. 1999. "Against Structural Universals". In *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
Linnebo, Ø. 2018. *Thin Objects: An Abstractionist Account*. Oxford: Oxford University Press.
McDowell, J. 1998. *Meaning, Knowledge, and Reality*. Cambridge: Harvard University Press.
Madl, T., Chen, K. Montaldi, D., and Trappl, R. 2015. "Computational Cognitive Models of Spatial Memory in Navigation Space: A Review." *Neural Networks* 65: pp. 18-43.
Ockham, W. 1957. *Summa Logicae*, in his *Philosophical Writings, A Selection*, ed. and trans. P. Boehner. London: Nelson.
O'Keefe, J., and Nadel, L. 1978. *The Hippocampus as a Cognitive Map*. Oxford: Clarendon University Press.
Parsons, C. 2008. *Mathematical Thought and its Objects*. Cambridge: Cambridge University Press.
Peacocke, C. 1992. *A Study of Concepts*. Cambridge: MIT Press.
Prinz J. 2011. "Has Mentalese Earned its Keep? On Jerry Fodor's *LOT 2*." *Mind* 120: pp. 485-501.
Putnam, H. 1975a. *Mathematics, Matter, and Method: Philosophical Papers, vol. 1*. Cambridge: Cambridge University Press.
---. 1975b. *Mind, Language, and Reality: Philosophical Papers, vol. 2*. Cambridge: Cambridge University Press.
Pylyshyn, Z. 1984. *Computation and Cognition*. Cambridge: MIT Press.
Quine, W. V. 1960. *Word and Object*. Cambridge: MIT Press.
---. 1969. *Ontological Relativity and Other Essays*. New York: Columbia University Press.

---. 1980. *From a Logical Point of View*, 2nd ed. Cambridge: Harvard University Press.
---. 1981. *Theories and Things*. Cambridge: Cambridge University Press.
Rescorla, M. 2009. "Cognitive Maps and the Language of Thought," *The British Journal for the Philosophy of Science* 60: pp. 377-407.
---. 2013. "Millikan on Honeybee Navigation and Communication." In *Millikan and Her Critics*, eds. D. Ryder, J. Kingsbury, and K. Williford. Malden: Wiley-Blackwell
---. 2015. "Bayesian Perceptual Psychology." In *The Oxford Handbook of the Philosophy of Perception*, ed. M. Matthen. Oxford: Oxford University Press.
---. 2016. "Bayesian Sensorimotor Psychology." *Mind and Language* 31: pp. 3-36.
---. 2017a. "From Ockham to Turing --- and Back Again." In *Turing 100: Philosophical Explorations of the Legacy of Alan Turing (Boston Studies in the Philosophy and History)*, eds. A. Bokulich and J. Floyd. Berlin: Springer.
---. 2017b. "Levels of Computational Explanation." In *Philosophy and Computing: Essays in Epistemology, Philosophy of Mind, Logic, and Ethics*, ed. T. Powers. Berlin: Springer
---. 2018. "Maps in the Head?", *The Routledge Handbook of Philosophy of Animal Minds*, eds. K. Andrews and J. Beck. New York: Routledge.
---. 2020a. "How Particular Is Perception?". Philosophy and Phenomenological Research 100: pp. 721-727.
---. 2020b. "A Realist Perspective on Bayesian Cognitive Science." In *Inference and Consciousness*, eds. A. New and T. Chan. New York: Routledge.
---. Forthcoming a. "Bayesian Modeling of the Mind: From Norms to Neurons." *WIREs Cognitive Science*.
---. Forthcoming b. "Perceptual Co-Reference." *Review of Philosophy and Psychology*.
Rips L. 1994. *The Psychology of Proof: Deductive Reasoning in Human Thinking*. Cambridge: MIT Press.
Rips, L., and Hespos, S. 2011. "Rebooting the Bootstrap Argument: Two Puzzles for Bootstrap Theories of Concept Development." *Behavioral and Brain Sciences* 34: pp. 135-136.
Schneider, S. 2011. *The Language of Thought: A New Philosophical Direction*. Cambridge: MIT Press.
Skinner, B. F. 1938. *The Behavior of Organisms*. New York: Appleton-Century-Crofts.
Soames, S. 2010. *What Is Meaning?*. Princeton: Princeton University Press.
Stalnaker, R. 1984. *Inquiry*. Cambridge: MIT Press.
Stich, S. 1983. *From Folk Psychology to Cognitive Science*. Cambridge: MIT Press.
Tolman, E. 1948. "Cognitive Maps in Rats and Men." *Psychological Review* 55: pp. 189-208.
Trommershäuser, J, Körding, K., and Landy, M., eds. 2011. *Sensory Cue Integration*. Oxford: Oxford University Press.
van Gelder, T. 1992. "What Might Cognition Be, If Not Computation?" *Journal of Philosophy* 92: pp. 345–381.
Wetzel, L. 2009. *Types and Tokens*. Cambridge: MIT Press.