

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Recognition is an Interface Between Perception and Cognition

Permalink

<https://escholarship.org/uc/item/9ws3t2bs>

Author

Abid, Greyson

Publication Date

2023

Peer reviewed|Thesis/dissertation

Recognition is an Interface Between Perception and Cognition

By

Greyson Abid

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Philosophy

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor John Campbell, Co-Chair

Professor Alva Noë, Co-Chair

Professor Michael Silver

Fall 2023

Abstract

Recognition is an Interface Between Perception and Cognition

by

Greyson Abid

Doctor of Philosophy in Philosophy

University of California, Berkeley

Professor John Campbell, Co-Chair

Professor Alva Noë, Co-Chair

My dissertation investigates a neglected topic within the philosophy of mind: our capacity to recognize people, places, and objects. Some examples of the phenomenon are vivid and infrequent, such as recognizing a tiger in the forest; other examples are mundane and pervasive, such as recognizing a colleague in the hallway. I defend the view that recognition is an interface between perception and cognition, serving as a point of interaction between the two domains. Without a capacity to recognize *types* of things, such as dogs or faces, our sense impressions would not be subsumed into our knowledge of general categories. We would not comprehend people, places, and objects as such on the basis of perception alone. Without a capacity to recognize *instances* of types, such as a neighbor's dog or friend's face, we would fail to integrate our sense impressions with our knowledge of particulars. Our situation would be like that of a Capgras patient, who knows they are in a relationship, appreciates that the person next to them perfectly resembles their partner, but does not register that the person *is* their partner.

Chapter 1 focuses on the question: Is recognition more like seeing, tasting, and smelling (perception) or judging, reasoning, and knowing (cognition)? On my view, it constitutively involves elements of both. It is partly perceptual because it exhibits a paradigmatically perceptual sensitivity to perspective, evidenced by our ability to discern subtle changes in viewpoint in object recognition tasks. It is partly cognitive because of its dependence on stimulus-independent representations housed in long-term memory—a claim I establish by examining empirical models of face recognition. I also introduce a framework for understanding the relation between familiarity, recollection, and the recognition of types and instances of types.

Chapter 2 inquires whether affect plays a constitutive role in recognition. According to *Constitutivists*, part of what it is to recognize a person, place, or thing is to enjoy a certain affective response. In contrast, *Causalists* hold that affect is only causally implicated in recognition. I discuss Capgras syndrome as a case study for exploring the disagreement between Causalists and Constitutivists. I suggest three ways of resolving the disagreement and show why none are particularly satisfying.

Chapter 3 considers an explanation of recognition in terms of a process of matching present and past impressions. This explanation mirrors subpersonal models of recognition I discuss in Chapter 1 but extends them to the personal level. I reject this explanation on the grounds that it is an instance of the broader mistake of trying to explain recognition in terms of a subject's more

basic mental capacities. While recognition constitutively involves both perception and cognition, it is nonetheless an *irreducible* mental capacity, in the sense that there are no mental capacities X such that the recognitional capacity fully depends on X . I show how we can establish this claim of irreducibility by treating it as a live empirical hypothesis. While this hypothesis is, strictly speaking, unverifiable, it is nevertheless supported by a large body of existing neuropsychological, ethological, and developmental evidence.

Chapter 4 applies my view to a failure of the recognitional capacity with significant practical consequences: the *cross-race effect*, traditionally characterized as our difficulty in recognizing individuals of other races. I argue that two distinct factors in fact comprise the effect: a difficulty in recognizing individuals of other races (Factor 1) and a corresponding lack of metacognitive awareness of this difficulty (Factor 2). Contemporary theories take Factor 1 as their sole *explanandum*. The existence of Factor 2—evidenced by our overconfidence in our ability to recognize individuals of other races—leaves these theories incomplete. I hypothesize that Factor 2 is a product of our tacit yet mistaken assumption that we recognize individuals of all races equally well. This hypothesis offers more accurate predictions than an alternative hypothesis which appeals to our tendency towards overconfidence as task performance worsens.

Acknowledgements

When I imagine what this dissertation would have been like without the help of colleagues, advisors, friends, and family, I imagine a pile of blank pages. Let me briefly acknowledge those who have helped me fill those pages in.

First, I thank both of my co-chairs John Campbell and Alva Noë for their support throughout graduate school. It is no secret that John and Alva are creative, brilliant, and insightful philosophers. Through countless discussions, they have widened and refined my philosophical eye, and pushed me to explore new ideas, perspectives, and opportunities. Perhaps more importantly, though, they have been there for me personally, treating me with kindness, patience, and care. I am deeply indebted to them both.

John and Alva were not my only advisors during my time in graduate school. I thank Michael Silver for welcoming me to his lab, serving as my outside committee member, and patiently walking me through new research involving visual attention. I also thank Ned Block, who supervised my semester-long visit to NYU. Ned has always been generous with his time and has fundamentally shaped my philosophical outlook.

So much of the joy of philosophy comes from philosophizing with others. During my time in graduate school, I've had the privilege of interacting, conversing, and learning from many amazing philosophers and philosophically inclined-scientists. Thank you to Dustin Stokes, Wayne Wu, Carlos Montemayor, Geoff Lee, Seth Yalcin, Susanna Siegel, Dave Chalmers, Rob Long, Jackson Kernion, Antonia Peacocke, Nico Orlandi, Tania Lombrozo, Eric Mandelbaum, Urte Laukaityte, Adam Bradley, Sam Clarke, Jessie Munton, Ravit Dotan, Blake Harper, Alberto Tassoni, Steven Gross, Chaz Firestone, Hakwan Lau, John MacFarlane, Kia Torab, Tyler Wilson, Mike Martin, Madeleine Levac, Adam Lovett, Emily Perry, Ezra Rubenstein, Rebecca Zhu, Mariel Goddu, Mathias Boehm, Jes Heppler, Sven Neth, Melissa Fusco, and many others.

I thank the students who I have had the privilege to teach and learn from during my time at Berkeley. Time together in the classroom was time well spent.

I thank my friends in Berkeley in other departments for keeping me sane (especially the psychologists and ecologists). You've reminded me that just about everything can be tied back to philosophy, but also that not everything should be. I thank my non-academic friends for all the outdoor trips together. Sometimes fly fishing is the best cure for writers' block.

I thank my partner Madeleine Welsch for her love and encouragement. Some of my fondest writing memories come from our time together in Vinalhaven, Maine. I am so grateful to have you by my side.

I thank my parents, who always encouraged me to ask questions, even when the answers weren't readily available. Both sacrificed so much so that I could pursue a high school, college, and graduate education and made sure I did not lose track along the way.

Let's hope there's nobody I forgot to recognize.

Chapter 1: Recognition and the perception-cognition divide¹

1. Introduction

Recent discussions in philosophy and psychology have focused on the distinction between perception and cognition.² This interest is not entirely new. Philosophers dating back to Aristotle have found the categories of perception and cognition to be theoretically fruitful ways of carving up the mind. Intuitively, the distinction is not difficult to appreciate. There is clearly *some* difference between seeing, touching, or tasting apple juice, on one hand, and thinking, reasoning, or making judgments about it, on the other. Intuition is partially vindicated by the success of scientific psychology, which readily employs such a distinction.

But how should recognition be understood in light of the distinction between perception and cognition? As a first pass (to be clarified shortly), the sense of recognition in question involves a sensitivity to particulars from one's past. Recognizing a familiar person (e.g., a colleague from work) is one instance of this, as is recognizing a place or thing that one has viewed before (e.g., a lake one visited as a child or one's jacket on a restaurant coat rack). It is not immediately apparent where recognition falls along the perception–cognition divide.

With a few notable exceptions, the topic of recognition (in the aforementioned sense) has been largely ignored in the philosophy of mind. This is surprising, given philosophers' interests in closely related matters, such as perceptual learning, imagination, and attention. One of the aims of this dissertation is to reignite philosophical interest in the topic. A natural starting point for any philosophical treatment of recognition involves determining what kind of mental phenomenon it is. Compare: if we discovered a new or long-forgotten chemical substance, a natural first step would be to determine its basic chemical kind.

A better understanding of recognition in light of the perception–cognition divide also helps clarify existing philosophical work on the significance of recognition. Strawson and Evans both emphasize the role of recognition in the identification of particulars. Strawson (1959) suggests that we would be unable to identify particulars at all if we could not recognize—or “reidentify”—them. Similarly, Evans (1982) argues that the capacity for recognition offers a way of identifying particulars that is more rudimentary than descriptive forms of identification. On Evans's view, recognition provides a means of identifying, for example, the Speaker of the House of Representatives *without* entertaining a corresponding description in thought. Apart from its role in identification, Peacocke (1992) highlights that the capacity for recognition is a prerequisite for a range of concepts, such as the concept “Lincoln Plaza,” whose possession requires a recognitionally-based way of thinking about their referents—what he calls *recognitionally concepts*. In considering the foundational role of recognition in identification and concept possession, it is natural to ask: What is needed to possess the capacity for recognition in the first place? Are only perceptual resources needed? Must cognitive resources be brought to bear? If so, which? By answering these questions, we more effectively spell out the implications of these works—an issue I return to in Section 7.

¹ The material in this chapter has been adapted from Abid (2021).

² See, for example, Block (2014), Firestone and Scholl (2016), Phillips (2017), Beck (2018), Mandelbaum (2018), and Green (forthcoming).

I begin by providing some clarificatory remarks concerning the perception-cognition distinction (Section 2) and recognition (Section 3). In Section 4, I offer a diagnostic tool for determining whether a mental phenomenon is at least partly perceptual. In Section 5, I use this diagnostic tool to make a case that recognition is at least partly perceptual. In Section 6, I draw on considerations relating to stimulus-independence and long-term memory to provide a qualified argument that recognition is also partly cognitive. The upshot is that recognition straddles the border between perception and cognition. I conclude in Section 7.

2. The perception-cognition distinction

A non-cursory discussion of the perception–cognition distinction would go beyond the scope of this chapter. Nevertheless, it is worth making a few preliminary remarks. First, the distinction is neither exhaustive nor exclusive. It is not exhaustive in that certain affective states, such as undirected forms of depression or boredom, may be neither perceptual nor cognitive. It is not exclusive in that there may be capacities that share features of both perception and cognition. In the case of “borderline” capacities, it is indeterminate whether the capacity in question is perceptual or cognitive, analogous to the way in which it is indeterminate whether a virus is living or non-living (Block, 2014). In the case of “interface” capacities, the capacity in question is a (determinate) hybrid that is part perceptual and part cognitive, analogous to the way in which a liger is a hybrid animal that is part lion and part tiger. Borderline and interface capacities are both compatible with a perception–cognition distinction. Indeed, I ultimately argue that recognition is an instance of the latter.

A hotly debated issue is whether top-down influences from cognition to perception impugn the hypothesis that perceptual systems (or “modules”) are informationally encapsulated from central cognition. (A system X is “informationally encapsulated” from a system Y if and only if X cannot perform computations over information stored in Y. Less formally, such a system X cannot use system Y as an informational resource.) Clark (2013) and Lupyan (2015) claim that such top-down influences threaten the very existence of a perception–cognition distinction. If these theorists are correct, an attempt to understand recognition in light of the perception–cognition divide is fundamentally misguided.

Is this threat serious? There are two possibilities to consider. The first is that the hypothesis that perceptual systems are informationally encapsulated from central cognition withstands empirical scrutiny, rendering the threat empty. Perhaps all existing evidence against the hypothesis is methodologically flawed (Firestone & Scholl, 2016). Perhaps top-down influences, such as those mediated by attention, can be understood in a way that does not threaten the hypothesis (Quilty-Dunn, forthcoming). The second possibility is that the hypothesis is falsified. This possibility threatens the perception–cognition distinction *only if* one grounds the distinction in a form of informational encapsulation; so, if the hypothesis turns out to be false, one option—apart from abandoning the distinction—would be to find an alternative ground for it, such as a difference in representational format.

In any case, while I do presuppose that there is a perception–cognition distinction, my aim is not to characterize what grounds it. One benefit of this modesty is that it allows us to remain open to a range of considerations in assessing where recognition falls along the perception–cognition divide.

3. Recognition

The lay concept of “recognition” is too ambiguous for philosophical discussion. In this section, I offer a more regimented characterization. After isolating the relevant sense of recognition (Section 3.1) and discussing its representational content (Section 3.2), I clarify how recognition (in the relevant sense) is distinguished from neighboring phenomena (Sections 3.3 to 3.4) and its relation to recollection and familiarity (Section 3.5), the personal-level (Section 3.6), and recognitional judgment (Section 3.7).

3.1 What is the relevant sense of recognition?

As noted earlier, my focus is on a sense of recognition that bears an intimate connection to one’s past. I call this token-recognition, which is to be distinguished from what I call type-recognition. “Type-recognition” occurs when a subject groups a particular into a certain category on the basis of its appearance.³ Without further qualification, this gloss is overly inclusive. A creature capable of color perception may group a sofa (that happens to be red) into the category of red things on the basis of its appearance, but few vision scientists would consider this is an instance of type-recognition. Similar considerations apply to certain other forms of low-level perception, such as depth or lightness perception. Why do we countenance face and object recognition, but not depth, color, or lightness recognition? A first pass answer is that in type-recognition a particular is grouped using a “sortal” category, roughly, a category expressed using a count noun (i.e., a quantifiable noun that can take both singular and plural forms) that specifies a thing’s (nominal or real) essence. This qualification is still not restrictive enough since type-recognition does not involve just any sortal category. Typically, the sortal categories are “basic-level” in Rosch’s (1978) sense (e.g., DOG, CAR), striking a balance between more specific subordinate categories (e.g., GREYHOUND, ACURA) and more general superordinate categories (e.g., ANIMAL, VEHICLE).⁴ A substantive question, which I do not discuss, is whether type-recognition is perceptual (see Mandelbaum, 2018).

Important for our purposes is that the type-recognition of some particular does not require any previous exposure to that very particular, although it arguably requires previous exposure to particulars belonging to the relevant category. For example, I have been exposed to plenty of dogs in the past, so I have no trouble type-recognizing my cousin’s newly adopted dog, even as it greets me for the very first time.

In contrast, “token-recognition” requires that an observed particular be registered as having been observed before. Here, “observation” should be construed broadly to include both direct and indirect observation. There are two reasons for this. First, an observed particular can be token-recognized even if it has not been *directly* observed before. For example, I might token-recognize Hugh Grant despite having never encountered him before in the flesh. Second, a previously observed particular can be token-recognized even if it is not *directly* re-observed. I might, say, token-recognize a friend in a painting. Nevertheless, in such cases there is *indirect* prior observation or re-observation through a representational medium. In the case of Hugh, the representational medium that allows for indirect prior observation is one of his films. In the case

³ Examples of particulars include people (better: individuals), places, and things.

⁴ I use capital letters to denote categories.

of my friend, the representational medium that allows for indirect re-observation is the painting that depicts them.

3.2 What is represented in token-recognition?

A token-recognized particular is *registered* as having been observed before. Talk of “registration” allows us to remain neutral regarding the difficult issue of what exactly is *represented* when one token-recognizes a particular. By way of example, consider the proposal that token-recognition involves representing an identity relation. On one version of this proposal, the content of such a representation might be expressed as “that₁ F = that₂ G”. Here, “that₁” denotes a singular element that functions to refer to a *currently* observed particular, “that₂” denotes a singular element that functions to refer to a *previously* observed particular (similar to a memory based demonstrative), and “F” and “G” denote general elements that function to attribute (possibly identical) general features to the particulars referred to by the respective singular elements.⁵

This proposal raises thorny questions. Why think identity is the represented relation—as opposed to a relation of resemblance? And, supposing identity is represented, must the representing subject possess the concept of identity—as opposed to some nonconceptual analog of identity? There is a complex interplay between the answers to these questions and the topic at hand, namely, where token-recognition falls along the perception-cognition divide. For instance, if perception is nonconceptual in the sense that a subject can perceive that *p* without possessing the concepts that characterize *p* (cf., Byrne, 2005) and token-recognition falls on the perception side of the perception–cognition divide, the representation of identity involved in token-recognition must not require the representing subject to possess the concept of identity.⁶ Even if we remain agnostic about issues concerning representation, it is worth remaining sensitive to these complexities.

3.3 Are type-recognition and token-recognition dissociable?

Clearly, type-recognition can occur without token-recognition. What about token-recognition without type-recognition? Certain examples are suggestive. I see a dog running through the trees at dusk without realizing it is a dog. I only discern that it is a medium-sized creature with brownish fur. Nevertheless, I am sure that it is the very same creature that was running through the trees a day ago. This is plausibly an instance of token-recognition without type-recognition. Of course, rejecting that token-recognition requires type-recognition is compatible with accepting that token-recognition is facilitated by type-recognition. I might have had an easier time token-recognizing the creature if I first grouped it into the category DOG. With these subtleties in mind, I tentatively accept that type-recognition and token-recognition are doubly dissociable and use “recognition” to refer exclusively to token-recognition in what follows.

⁵ This approach extends Burge’s (2010) framework.

⁶ One might try to bypass these issues by sketching an alternative proposal. Perhaps a token-recognized particular is represented in a self-referential fashion using just one singular element that functions to refer to a currently observed particular. Building on the above proposal, the content of this representation might be symbolized as “that F which I have observed before”. On this alternative, the previously observed particular is represented *as such*. Nevertheless, this proposal raises its own thorny questions of whether the representing subject must possess indexical concepts or the concept of observation. There is again a complex interplay between the answers to these questions and the issue of where token-recognition falls along the perception–cognition divide.

3.4 How is recognition different from short-term reidentification?

Recognition is distinct from various forms of short-term reidentification. Consider watching a cyclist weave through the streets. The cyclist may not always be in your direct line of sight as you track them. They may momentarily disappear behind a truck. Still, you continually track the cyclist through these brief occlusions. This phenomenon has received a great deal of attention among vision scientists. It is now well-established that visually tracking a particular through occlusion is subject to demanding spatiotemporal constraints. For example, visually tracking an object through occlusion is significantly impaired if the distance between the position of disappearance and reappearance of the object under occlusion is too large or if a small temporal lag is introduced during the occlusion period (Flombaum & Scholl, 2006). While our ability to recognize a particular certainly degrades over time, recognition is not subject to such highly circumscribed spatiotemporal constraints. Recognition is “spatiotemporally robust” in that a particular can be recognized even if it was last observed many months or even years before, perchance in a distant location. Just consider the possibility of recognizing a long-lost childhood friend while travelling in foreign country.⁷

Similar considerations apply to other forms of short-term reidentification. If I am shown an image of a face and am told that I will have to reidentify it a few seconds later, I may try to actively retain the image in visual working memory to complete the task. This strategy is only effective over small timescales since information held in visual working memory degrades within ~10 seconds (Zhang & Luck, 2009; Ricker & Cowan, 2010). Recognition is not subject to this constraint.

The claim that recognition is spatiotemporally robust is to be understood as a psychological generalization. Psychological generalizations (in general) are not exceptionless and usually contain hidden *ceteris paribus* clauses. It is for this reason that an intoxicated individual who can only hold one item in working memory is no counterexample to the psychological generalization that working memory has a capacity of ~4 items. The same holds true for many purported counterexamples to the generalization that recognition is spatiotemporally robust.

3.5 How does recognition relate to familiarity and recollection?

Registering that an observed particular has been observed before—in a manner not subject to the sorts of highly circumscribed spatiotemporal constraints mentioned above—*suffices* for recognition. More specifically, it suffices for a type of recognition known as “familiarity”. Familiarity *does not* require the retrieval of contextual information from past episodes involving that particular, such as when or where one saw it. Walking down the street, I might cross paths with a familiar person I have met years before without being able to retrieve specific information about where or when I met them. James has an example of familiarity with a *postponed* retrieval of contextual information (or “recollection”) in mind in the following passage:

...I enter a friend’s room and see on the wall a painting. At first I have the strange, wondering consciousness, “surely I have seen that before”, but when or how does not become clear. There only clings to the picture a sort of penumbra of familiarity—when

⁷ Larzabal et al. (2018) provide an illustration of the spatiotemporal robustness of recognition by showing that individuals recognize pictures last observed 8-14 years earlier at a rate above chance.

suddenly I exclaim: “I have it, it is a copy of part of one of the Fra Angelicos in the Florentine Academy—I recollect it there!” (James, 1890; cited in Yonelinas, 2010).

According to now widely accepted “dual-process” theories (Mandler, 2008; Yonelinas, 2010), separate processes underlie familiarity and recollection. The former involves a gradable signal detection process based in the perirhinal cortex. The latter involves an all-or-nothing memory retrieval process based in the hippocampus. Dual-process theorists hold that reidentification tasks *often* involve both underlying processes.

3.6 Is recognition personal or subpersonal?

On the views of Strawson, Evans, and Peacocke, it is a *subject’s* capacity for recognition that is crucial for identification and concept possession. More generally, we attribute the capacity for recognition to a subject as a whole, not to neural or information-processing systems within that subject.⁸ Of course, there may be subpersonal-level *correlates* of recognition. For instance, exposure to familiar stimuli result in reduced neural responses. This is known as “repetition suppression”. Repetition suppression occurs in early visual areas, such as V2 (Huang et al., 2018). Nevertheless, V2 does not engage in “recognition” in any non-figurative sense, and I claim it would be a category mistake to suggest otherwise. In Section 6, I discuss a model of recognition that employs subpersonal-level representations and operations on those representations.

Given this personal-level characterization, how are we to understand devices such as facial recognition systems? We might claim that these devices are only capable of recognition in a figurative sense—much like we claim that a thermostat only figuratively “knows” the temperature. Or, we might claim that the personal-level extends to such devices—much like we often extend the personal-level to non-human animals not generally regarded as persons. This, in turn, would allow us to hold that these devices are literally capable of recognition. Our intuition sits somewhere between these two extremes, varying depending on the configuration and behavior of the device in question along with contextual factors that raise or lower our standards for mentalistic attributions.

3.7 What is the difference between recognition and judgment?

It is crucial to distinguish recognition from recognitional judgment. A “judgment” is a doxastic attitude, similar to belief, that a subject takes towards some proposition *p*. Unlike non-doxastic attitudes (e.g., desiring), doxastic attitudes are governed by norms describing the conditions for epistemic justification and permissibility—what are often called “epistemic” norms. An example: If one is inclined to judge that *p* but discovers that the evidence for *p* is equivocal, all else equal, one ought to suspend judgment about *p*. This judgment would be epistemically impermissible.

A “recognitional” judgment is a special instance of a judgment—one made when an observed particular is judged as having been observed before. In contrast, recognition does not require one to form any such judgement.

There are a few ways that recognition and recognitional judgment come apart. I might recognize a person on the street and only *later*, in an act of self-reflection, judge that I have

⁸ For a discussion of the personal-level versus subpersonal-level distinction, see Drayson (2014).

seen them before. Here, my recognition of the person acts as a (partial) causal antecedent for my recognitional judgement; it also plays a normative role in justifying the recognitional judgment. Alternatively, I may see a stranger on the street and mistakenly judge that I have seen them before. In this case, successful recognition does not serve as a causal or justificatory precursor to my recognitional judgment. Finally, it is plausible that many instances of recognition occur without any recognitional judgment at all. Russell (1921) offers the example of a horse returning to its stable as one such instance. More generally, it is plausible that some animals and infants have the capacity for recognition but not recognitional judgment.

3.8 Recap

To summarize, the relevant sense of recognition is token-recognition, which occurs when one registers—in a spatiotemporally robust manner—that an observed particular has been observed before. Recognition is a personal-level capacity that does not require the retrieval of contextual information from past episodes involving that particular. Finally, recognition is distinct from recognitional judgment.

4. The perspectival test

Upon encountering a particular, we often exhibit a response that varies across certain changes in our perspectival relations to that particular. I call this type of variable response pattern *perspectival sensitivity*. With a few clarifications (Section 4.1 and Section 4.3), I argue that a mental phenomenon is at least partly perceptual if it exhibits perspectival sensitivity.⁹ My argument is inductive: Paradigmatically perceptual phenomena exhibit perspectival sensitivity (Section 4.2), whereas cognitive phenomena do not, at least not independently of perception (Section 4.4).

4.1 What exactly is a perspectival relation?

A non-exhaustive list of perspectival relations between some subject *S* and some particular *o* includes: The spatial position of *o* with respect to *S* along with *o*'s surrounding context and various conditions of the environment (e.g., illumination and background noise conditions). An exhaustive list would include all and only those relational features that determine *S*'s *perceptual* perspective on *o*—what we often speak of as the viewing conditions (of *o* with respect to *S*). *S*'s perceptual perspective on *o* is to be distinguished from *S*'s *cognitive* perspective on *o*—the sort of perspective *S* has while entertaining a thought about *o* under some specific mode of presentation. Putting things this way makes it apparent that there is an element of circularity in appealing to perspectival relations—and, by extension, perspectival sensitivity—to determine whether some mental phenomenon is perceptual: Understanding these notions requires some antecedent grip on the perception–cognition distinction. For our purposes, this circularity is relatively unproblematic. Problems *would* arise if we were to use these notions to articulate the distinction between a perceptual and a cognitive perspective or the grounds for the perception–cognition distinction.

4.2 Perspectival sensitivity in perception

⁹ Throughout, I have a broad extension of “phenomenon” in mind that includes states, events, processes, and capacities.

Let us consider some paradigmatically perceptual examples of perspectival sensitivity in which one's *subjective* response varies across certain changes in one's perspectival relations to a particular.¹⁰ Consider first the example of viewing two same-sized trees from different distances (e.g., 50 and 100ft.). It is often said that this example illustrates the dual aspect of perception. One aspect corresponds to the "apparent" sense in which the nearby tree appears larger than the more distant tree. The other corresponds to the "objective" sense in which the two trees appear to be same size.¹¹

The trees appear to be the same objective size because of the *constancy* of visual perception. *Ceteris paribus*, each tree appears stable in objective size across certain changes in the perspectival relations between the viewer and the tree, notably, certain changes in distance. (For brevity, I omit the *ceteris paribus* clause in what follows.) This is aptly called "size constancy". At the same time, the trees differ in apparent size because of the *perspectival sensitivity* of visual perception. The apparent size of each tree varies as one changes one's perspectival relations to the tree, specifically, one's distance from the tree. Thus, the tree viewed from 50ft. has a larger apparent size than the same-sized tree viewed from 100ft., occupying a greater portion of one's visual field.

Another example that demonstrates the constancy and perspectival sensitivity of visual perception involves color appearance.¹² Seeing a monochrome pink wall partially covered in shadows reveals both a uniform and variable color appearance. The uniform color appearance of the wall is made possible by the constancy of visual perception. This appearance stays constant across certain changes in one's perspectival relations to different portions of the wall, changes which alter viewing position, color context, illumination, and so on. This is called "color constancy". In contrast, the variable color appearance of the wall is made possible by the perspectival sensitivity of visual perception. This appearance does not stay constant across the aforementioned changes in perspectival relations to different portions of the wall; analogous to apparent size, variable color appearance is sensitive to one's perspectival relations to different portions of the wall.

The constancy and perspectival sensitivity of visual perception extends to other sense modalities as well. As one moves closer to a concert venue's speaker system, there is a clear sense in which the music seems to get louder, as evidenced by the fact that, for example, I must shout to be heard when near the speaker. This is compatible with it seeming to me that the actual *source volume* remains roughly the same. This happens because we can distinguish alterations in the perceived loudness of a sound of constant intensity. The perceived loudness of the sound is made possible by the perspectival sensitivity of audition. Perceived loudness varies across certain changes in distance to the source (e.g., the speaker system). The perceived intensity of the sound is made possible by the constancy of auditory perception. The perceived intensity of the sound

¹⁰ In Section 5, I focus on examples of perspectival sensitivity involving *behavioral* responses.

¹¹ This example is given by Peacocke (1983), Tye (2000), Noë (2004), Schellenberg (2008), and Lande (2018). Schwitzgebel (2006) is skeptical as to whether perception possesses a dual aspect. Although disputes concerning phenomenology are always difficult to resolve, there is ample evidence that lay perceivers can readily distinguish "apparent" from "objective" appearances, suggesting that the dual aspect of perception is not merely a philosophical construct (Wagner, 2006; Green & Schellenberg, 2018; see also Morales et al., 2020).

¹² This example is given by Noë (2004) and Schellenberg (2008).

does not vary across changes in distance to the source, at least within a certain range. This is somewhat confusingly called “loudness constancy” (Zahorik & Wightman, 2001).

4.3 The perspectival test

I contend that we can employ perspectival sensitivity as a diagnostic tool. According to what I call the “perspectival test”, *a mental phenomenon is at least partly perceptual if it exhibits perspectival sensitivity*.

Three clarifications are in order. First, the perspectival test indicates a mental phenomenon is at least partly perceptual *in the absence of defeaters*. For example, olfactory states exhibit perspectival sensitivity. Burge, however, does not regard olfactory states as perceptual on the grounds that they are not *about* (in the intentional sense) distal features of the environment.¹³ While I am skeptical of Burge’s analysis here, I note that it does not conflict with the perspectival test since a failure to exhibit the requisite form of intentionality required for perception would qualify as a defeater.

Second, the perspectival test is modest. It only indicates that a mental phenomenon is *at least partly* perceptual. It does not specify whether the phenomenon in question is purely or only partly perceptual. What does it mean to be *partly* perceptual? It cannot simply amount to being causally influenced by perception; this condition is far too weak to be of any theoretical interest. Rather, a mental phenomenon is “partly perceptual” just in case it in part constitutively depends on some perceptual element (e.g., a perceptual mechanism or representation). “Partly cognitive” is defined similarly. In Section 6, I argue that recognition is partly cognitive in virtue of its partial constitutive dependence on long-term memory representations. This claim does not conflict with the perspectival test. Rather, it implies that recognition is partly perceptual and partly cognitive.

Third, the perspectival test does not require that all perceptual phenomena exhibit perspectival sensitivity. Hallucinations are arguably perceptual, at least according to non-disjunctivists who claim that hallucinations and non-hallucinatory perceptual experiences belong to the same fundamental kind. Yet, hallucinations cannot always exhibit perspectival sensitivity since we are not perceptually related to any particulars in instances of “pure” hallucination. This is not in tension with the test. Passing the test is sufficient but not necessary for counting as at least partly perceptual.

Similar considerations apply to certain forms of imagination. Suppose I see a tree from 50ft away, turn so I no longer see it, start to imagine it, and move 100ft away from it—in that order. As I imagine the tree, its apparent size does not vary across *actual* changes in my distance from it, though it may vary across *imagined* changes in distance. While this form of imagination fails to exhibit perspectival sensitivity, this does not preclude the possibility that it is at least partly perceptual—a conclusion we may reach in light of certain neural, phenomenological, or representational similarities between imagination and perception.

¹³ Burge regards olfaction as non-perceptual because it does not result in “objectification.” According to Burge, “*Objectification* is [the] formation of a state with a representational content that is *as of* a subject matter beyond idiosyncratic, proximal, or subjective features of the individual” (2010, p. 397). For discussion suggesting that olfaction does result in objectification, see Begby (2011).

4.4 Assessing the adequacy of the perspectival test

Do cognitive phenomena ever exhibit perspectival sensitivity? Of course, there are cases in which we form *perceptually-based* cognitive states that *reflect* the perspectival sensitivity of perception. A painter, for instance, might make a judgement about the apparent size of a tree while attempting to depict it. Other examples might include certain perceptual demonstrative thoughts, thoughts about particulars—standardly expressed through demonstratives, such as “this” or “here”—made available by one’s perceptual relations to those particulars. Yet, the fact that some perceptual feature, such as perspectival sensitivity, can be reflected in a perceptually-based cognitive state is no reason to hold that this feature is not a mark of the perceptual. After all, some features that are *clearly* perceptual, such as the apparent length of the lines in the Müller-Lyer illusion, can be reflected at the level of perceptual judgment. In the case of the Müller-Lyer, I may form the perceptually-based judgement that the lines look to be different lengths, even though I know they are not. The mere fact that I can form such a judgment does not show that the illusion has a non-perceptual basis. If anything, the fact that I cannot revise this perceptually-based judgement in light of my background knowledge is evidence that the illusion is a product of perception.

To assess the adequacy of the perspectival test, we should instead ask whether cognitive phenomena exhibit perspectival sensitivity *independently* of the perspectival sensitivity of perception. An affirmative answer is suggested by the fact that we often change our judgements and beliefs about things by literally changing our perspectival relations to them. For instance, I may believe that my bicycle is in the backyard but revise my belief if I change my perspectival relation to my bicycle by stepping outside. In turn, I may find that my bicycle was in my backyard *yesterday* but is now nowhere to be found. Notice, however, that this revised belief is not the result of a sensitivity to my perspectival relation to the bicycle *per se*. Rather, the change in my perspectival relation to the bicycle only results in a revised belief insofar as the change in relation yields new evidence. The change in perspectival relation plays only an instrumental role in changing my belief. I would not revise my belief if I were to step outside and find my bicycle in the same place as before. Likewise, I may have revised my belief in just the same way even if I did not step outside (e.g., if I were told that every bicycle in the neighborhood was stolen).

Similar considerations help to disarm other potential counterexamples to the perspectival test. Consider the game “hot and cold”, in which a small household object is hidden by player 1 while player 2’s eyes are closed. After the object is hidden, player 2 begins to search for it. As player 2 gets closer to the object, player 1 says, “hotter”. As player 2 gets farther away from it, player 1 says, “colder”. Now, player 2 may come to form the belief that they are *such-and-such* a distance away from the object, and this belief may be *evidentially* sensitive to the testimony of player 1, but the belief in question is not sensitive to player 1’s perspectival relations to the object *per se*. Again, this is easy enough to appreciate: Had player 3 come along and secretly moved the object mid-game, player 2 would have still revised their belief in accordance with player 1’s instructions.

5. The perspectival sensitivity of recognition

In this section, I provide evidence that recognition exhibits perspectival sensitivity. In conjunction with the perspectival test, these considerations establish that recognition is at least partly perceptual.

By way of introduction, consider an experiment by Brady et al. (2008), in which subjects were instructed to remember 2,500 object images from a range of different object categories over the course of 5.5 hours. Each image was displayed once for three seconds. An additional task that required subjects to note image repetitions was also included so that subjects paid attention to each image. After this initial display, subjects were presented with two images, one “novel” and one “old” (i.e., an image selected randomly from the 2,500 object images displayed earlier). Subjects performed a task requiring them to select the image they had seen in the initial display of the object images. (This type of task is known as a “two-alternative forced choice” task.) In the “novel” condition, the novel image was of an object that differed in category from any of the objects depicted earlier. In the “exemplar” condition, the novel image was of a physically distinct object of the same category as the object depicted in the old image. In the “state” condition, the novel image was of the same object depicted in the old image but was displayed in a different state or pose. Changes in the “state” of an object are identity-preserving changes in its initial condition. For example, a briefcase undergoes a state change in this sense if it goes from being unopened to being opened. Similarly, a dresser would undergo a state change if one of its drawers were opened. Changes in the “pose” of an object are changes in the perspective or vantage point from which it was initially depicted. Across all three conditions, subjects were remarkably accurate: They correctly identified the old image 93% of the time in the novel condition, 88% of the time in the exemplar condition, and 87% of the time in the state condition.

These results are interesting for many reasons.¹⁴ The main point which I wish to extract from this study, however, are the results from the state condition. The results of this condition indicate that our capacity for recognition goes beyond the capacity to distinguish previously observed objects within or across categories. In the state condition, subjects were significantly above chance in distinguishing changes in the state or pose of a single object. The finding that subjects can distinguish changes in the pose of an object *would* provide evidence that in recognizing a particular subjects’ responses vary across certain changes in their perspectival relations.

In Section 3, I claimed that registering that a particular (e.g., an object) has been observed before is necessary for recognition. Are participants in this experiment really registering that the objects depicted in the images in the forced choice task are the same objects depicted in the images they observed several hours earlier? Perhaps participants had a *raw feel* for which object image in the forced choice task was the correct one but were unable to *recognize* the depicted object. There is reason to think this worry is misplaced. As Brady and colleagues note: “Participant reports

¹⁴ For one, they provide a healthy contrast to change and inattention blindness studies. Such studies are frequently taken to show that we do not internally store rich worldly detail. For instance, Noë (2002) suggests that “work on change blindness seems to suggest that we may not in fact actually produce ... detailed internal models” (p. 9). As Noë notes, change blindness studies concern unattended aspects of a scene. Granting that change blindness studies do not merely reflect difficulties with a post-perceptual comparison process, it may be reasonable to infer from these studies that we do not generate rich, internal models of those unattended aspects (cf., Simons & Rensink, 2005). By forcing subjects to allocate their attention appropriately, Brady and colleagues provide evidence that we retain richly detailed information concerning previously *attended* aspects of a scene. These results are likely even more striking in everyday contexts, where objects and scenes tend to be viewed multiple times and for durations extending beyond just a few seconds. In a task requiring subjects to answer questions concerning the details of previously viewed natural scene images, performance greatly improved as scene exposure times increased and when the scene was viewed multiple times before the questions were asked (Melcher, 2006).

afterward indicated that they were usually explicitly aware of which item they had seen [several hours before], as they expressed confidence in their performance and volunteered information about the details that enabled them to pick the correct items” (2008, p. 14327).

Still, there is an alternative explanation of the participants’ results. In the state condition, each trial involves either (i) a change in the initial condition of the depicted object (its “state”), (ii) a change in the vantage point from which the object was initially depicted (its “pose”), or (iii) a change in the object’s state *and* pose. Given the inclusion of (i) and (iii), it is difficult to rule out the interpretation that subjects’ strong performance in the state condition is primarily driven by their capacity to distinguish changes in an object’s state, not its pose. If this were the case, one might argue that the study by Brady and colleagues does not provide strong evidence that recognition exhibits perspectival sensitivity. Instead, participants might only be sensitive to non-relational changes in the depicted object.

A discrimination task by Hollingworth and Henderson (2002) provides further evidence that recognition exhibits perspectival sensitivity while avoiding this latter worry. Using an eye tracker, the experimenters allowed subjects to freely view an image of a scene filled with various objects until they had directly fixated on a target object (e.g., a toy truck or a notepad). Once subjects had shifted their gaze to another part of the screen, the target object was covered with a speckled mask. After re-fixating the mask and pressing a button to indicate their readiness, subjects performed a two-alternative forced choice task in which they were to identify the original scene among two different versions displayed successively: The original version of the scene and a version of the scene that was indistinguishable from the original with the exception that the target object was rotated in depth by 90°. ¹⁵ (The order was counterbalanced across trials.) If subjects were *insensitive* to their perspectival relations to the target object in performing this recognitional task, their performance should be approximately at chance. Yet, participants were quite successful in recognizing changes in the orientation of the target object: They correctly identified the initial orientation of the target object well over 80% of the time. Similar results obtained in another experiment which involved much lengthier delays between the initial scene display and the discrimination task (i.e., 5–30 minutes) (2002, experiment 2).

One might worry that these studies do not show that subjects are sensitive to changes in their perspectival relations to particulars *per se*. Instead, subjects may only be sensitive to a particular’s surface features whose visibility is affected by changes in perspectival relations. For example, certain changes to the orientation of a hammer might render more of its handle visible while simultaneously rendering less of its head visible. Maybe a sensitivity to differences in these visible surface features alone explains the results of the above studies. If this is correct, these results do not establish the perspectival sensitivity of recognition.

A study by Standing et al. (1970) provides evidence in favor of the perspectival sensitivity of recognition while avoiding this worry. Using a two-alternative forced choice task, Standing and colleagues showed that subjects are capable of recognizing left–right mirror reversals in images. Even after significant 30-minute delays, subjects detected mirror reversals well over 80% of the

¹⁵ It is unlikely that a rotation *in depth* would make a target object consistent with typical orientations in the real world appear atypical or vice versa. For example, a toy truck sitting wheels down on a table would not suddenly appear atypical after a rotation in depth. By contrast, reorienting the toy truck so that its wheels face the air would render its appearance atypical.

time. These results are significant because mirror reversals alter one’s perspectival relations to particular without affecting which of its surface features are visible. This indicates that the perspectival sensitivity of recognition cannot be “explained away” in terms of a mere sensitivity to differences in visible surface features.

Let me mention a final concern. The issue of whether recognition exhibits perspectival sensitivity is closely connected to debates concerning the effects of viewpoint on object recognition. Translated into the context of these debates, the foregoing results would seem to indicate that object recognition is affected by changes in an object’s viewpoint (relative to the viewpoint from which it was initially observed). In general, however, the evidence that has emerged from these debates is mixed, and some studies suggest no effect of viewpoint on object recognition (see Hummel, 2013). Is this cause for alarm?

The answer would seem to be yes if the debate over the effects of viewpoint on object recognition is cast in “either/or” terms. Yet, this is probably a false dichotomy. Tarr and Hayward (2017) provide evidence that, by default, subjects concurrently encode *both* viewpoint-dependent and viewpoint-invariant information for the purposes of object recognition but flexibly produce viewpoint-dependent or viewpoint-invariant responses in object recognition tasks, depending on the nature of the task (see also Leek & Johnston, 2006). Similarly, these findings mesh with evidence that two neurally-dissociable subsystems underlie object recognition—a “viewpoint-dependent” and a “viewpoint-invariant” system—whereby certain contextual factors lead to one system dominating the other (Burgund & Marsolek, 2000). If a story along these lines is correct, we should *expect* studies to produce mixed results regarding the effect of viewpoint on object recognition.

Translated back into our larger discussion, these results suggest that object recognition exhibits *both* perspectival sensitivity and insensitivity—just as perception does. While I have focused on perspectival sensitivity in this section, the perspectival insensitivity of recognition would explain why subjects’ performance on recognition tasks is not typically impaired by changes in vantage point, lighting conditions, etc.

6. Recognition is an interface capacity

The last two sections indicate that recognition is at least partly perceptual. However, the perspectival test used to establish this conclusion leaves open whether recognition is purely or only partly perceptual. In this section, I explore two lines of thought in support of the hypothesis that recognition is an interface capacity that is partly perceptual and partly cognitive. The first concerns the role of background information in recognition (Section 6.1). The second concerns the role of long-term memory (Section 6.2). To foreshadow, I find support from the second line of thought—with one qualification.

I advance our discussion by considering a model of recognition that has remained popular for several decades. I have in mind Bruce and Young’s (1986) influential model of face recognition. In what follows, I omit details of the model that are not relevant for our purposes and focus on the model itself, rather than evidence in favor of it.

Given that our discussion is cast in terms of Bruce and Young’s model, a couple points are worth mentioning. First, the model is limited in scope and applies only to familiar face recognition.

This restriction is not entirely artificial since recognition has been researched extensively in the context of faces. Moreover, it will become apparent that our conclusion straightforwardly generalizes beyond face recognition. Second, it is always possible that new findings will undermine the model. Of course, this is true of any model. That said, the model has largely withstood the test of time. As Burton and colleagues remark: “In the 25 years, since Bruce and Young published their paper, the core theoretical distinctions it draws have survived remarkably intact” (2011, p. 953–954).

On Bruce and Young’s model, face recognition begins with a process of “structural encoding” in which the visual system produces a structural code of a face. A “structural code” is a representation of a face that abstracts away from certain transient features, such as hairstyle, while encoding certain invariant features of the face, such as the spatial arrangement of facial features. While some structural codes, such as those used for the analysis of facial speech, retain information concerning facial expression, those used for face recognition do not. They do, however, encode information concerning viewpoint.

These structural codes are transferred to “face recognition units” (FRUs), which store structural codes corresponding to previously observed faces—abstract long-term memory representations of previously observed faces. Each FRU yields an activation signal, which is a function of the degree of resemblance between the structural codes produced by the process of structural encoding and the structural codes stored in the FRU. A strong FRU activation signal is an indication that a face is familiar.

A FRU’s baseline activation signal can also be increased by the activation of “person identity nodes” (PINs). Unlike FRUs, PINs are not modality-specific. For instance, a PIN can be directly activated by a voice recognition unit—the auditory analogs of FRUs. In turn, PINs can activate and be activated by other person-specific memories held in central cognition, such as memories about where or when some individual was born. It is through the activation of PINs that top-down influences are exerted on FRUs.

6.1 The role of background information in recognition

Let us turn to a first line of thought suggesting that recognition is an interface capacity. One might think that recognition cannot be purely perceptual because it requires a kind of decision, namely, a decision that the particular that is currently being observed is the same particular that has been previously observed. Since Helmholtz, it has been widely claimed that perception requires a (possibly figurative) form of decision-making or inference. For instance, it is often said that perceptual systems rely on an assumption that light comes from above to infer the three-dimensional shape of an object from two-dimensional patterns of light projected on the retina. Nevertheless, one might suspect that the kind of decision-making needed for recognition requires access to a sophisticated web of background information. One reason for thinking this is that many of the particulars we observe and later recognize do not remain constant in appearance over time. And, if recognition is a task that requires unrestricted access to background information, one may worry that recognition is too computationally demanding a task for perceptual systems to solve without the aid of cognition.

On Bruce and Young’s model, the kind of “decision” required for recognizing a face is rudimentary: A FRU’s activation signal, *ceteris paribus*, depends on only: (i) the degree of

resemblance between structural codes and (ii) the activation of certain PINs. Still, one might wonder whether such a model is too simple to account for the complexity of recognizing a face. We often see faces from distinct vantage points, under different lighting conditions, with varying expressions, and so forth. How is the model supposed to account for these complexities?

From the get-go, Bruce and Young were sensitive to these issues, speculating that structural codes encode only invariant features of faces that are stable over time. A similar approach is endorsed by Burton et al. (2005), who hypothesize that the stored representations that facilitate face recognition are produced by a process of averaging that preserves information concerning invariant facial features while eliminating information concerning transient facial features. Consequently, superficial variations in facial appearance over time will not be disruptive. Longmore et al. (2008) offer a different approach, hypothesizing that face recognition is made possible by storing a collection of richly detailed representations, each produced as a result of a previous encounter with the face. As the collection corresponding to a particular face *F* grows, so does the likelihood that a novel representation of *F* will resemble one of the representations of *F* stored in the collection. Redfern and Benton (2019) suggest combining these two approaches.

Regardless, the systems responsible for face recognition account for variations in the appearance of a particular face over time *without* direct access to arbitrary background information. Face recognition is a complex task, but one that is completed using only face-specific information.

Consider the following rejoinder. The systems responsible for face recognition may not be informationally encapsulated from central cognition. After all, Bruce and Young's model leaves open the possibility that central cognition can influence the activation signals of FRUs, at least indirectly through the activation of PINs. If one holds that perceptual systems are informationally encapsulated from central cognition, one may worry that face recognition cannot be purely perceptual.

Care must be taken to distinguish the well-established claim that perceptual systems are not *causally* encapsulated from central cognition from the contentious claim that perceptual systems are not *informationally* encapsulated from central cognition. As Wu stresses, in discussions of modularity of mind:

The issue is not merely *causal* but *informational* encapsulation... reference to information is reference to semantic content over which computations are performed ... establishing the failure of encapsulation [of system X from system Y] requires providing a mechanism where... computations [of system X] have access to and use [system] *Y* as an informational content resource (Wu, 2013, p. 656).

On Bruce and Young's model, FRUs cannot perform computations over information—that is, semantic content—stored in central cognition. This is because no information is transferred from central cognition to PINs to FRUs. For example, the information *Marcus is from Seattle* stored in central cognition is not transferred to PINs, even if the retrieval of this information activates a certain PIN. PINs do not store biographical information. Since interactions from central cognition to FRUs are mediated by PINs, the information is not transferred to FRUs either. The systems responsible for face recognition are thus informationally (but not causally) encapsulated from

central cognition, according to Bruce and Young's model, because they cannot access information housed in central cognition, not even indirectly.¹⁶

6.2 The role of long-term memory in recognition

Let us consider another line of thought suggesting that recognition is an interface capacity. Long-term memory representations play a constitutive role in recognition. On Bruce and Young's model, they take the form of structural codes stored within FRUs. These representations are "stimulus-independent" in the sense that they are not causally sustained by a distal stimulus by means of a present proximal stimulus, such as a pattern of light impinging on the retina (Beck, 2018). Yet, it is often said that, unlike cognitive representations, perceptual representations are stimulus-dependent, suggesting that the structural codes in Bruce and Young's model fall on the cognition side of the perception-cognition divide.

In assessing this line of thought, it is important to distinguish a FRU's activation signal from the structural codes stored within a FRU. In cases of successful face recognition, a face is a distal stimulus that causally sustains the activation signal of the appropriate FRU by means of a proximal stimulus. So, unlike structural codes, a FRU's activation signal is stimulus-dependent.¹⁷ At the personal-level, these activation signals correspond to a stimulus-dependent recognitional response.

It might be argued that what distinguishes perceptual from cognitive representations is not stimulus-dependence *per se*, but that the former *function*, or aim, to represent in a stimulus-dependent fashion (cf., Phillips, 2017; Beck, 2018). One motivation for this amendment is that it allows certain representations to count as perceptual even in cases of hallucination; in such cases, a representation cannot be stimulus-dependent, though it might function to represent in a stimulus-dependent fashion. This amendment does not reduce the force of the present line of thought. It is not merely that long-term memory representations *are* stimulus-independent. It is apparent from reflecting on their role in our cognitive economy that they also function to represent in a fashion that is causally independent of any present (as opposed to past) stimulus. Thus, a long-term memory representation can be said to fulfil its function even if the represented stimulus no longer exists.

It might also be argued that certain perceptual representations nonetheless function to represent in a stimulus-independent fashion. The most compelling examples that come to mind

¹⁶ The distinction between causal and informational encapsulation adds a layer of interpretive complexity to studies which emphasize the role of conceptual processing in recognition. Schwartz and Yovel (2019) show that associating conceptual information with a face during an encoding phase facilitates later recognition. The experimenters rule out a range of explanations of this benefit (e.g., that it is due to more elaborate or global face processing) to show that it arises because of the formation of a conceptual representation. Nevertheless, the experimenters do not specify *how* this conceptual representation improves later recognition. On one interpretation, conceptually represented information improves recognition through an indirect causal influence on FRUs. On another interpretation, the systems responsible for face recognition perform computations over the conceptually represented information associated with the previously encoded face. Only the latter interpretation threatens the hypothesis that facial recognition systems are informationally encapsulation from central cognition. Yet, it is unclear why this interpretation ought to be favored over the former.

¹⁷ To claim that a FRU's activation signal is stimulus-dependent is not to deny that there are other factors relevant to sustaining and altering it.

involve so-called “assumptions” made by perceptual systems, such as the assumption made by the visual system that light tends to come from above. While new experiences may cause one to revise such assumptions (Adams et al., 2004), arguably, the aim of these assumptions is to represent the environment in a way that is not causally sustained by any present proximal stimulus.

Yet, a popular view is that these assumptions are not explicitly represented by perceptual systems at all. They are what Pylyshyn (1999) calls “natural constraints”. As he puts it, “[t]he visual system does not need to access an explicit encoding of the constraint: it simply does what it is wired to do, which, as it happens, means that it works in accordance with the constraint discovered by the theorist” (p. 354; see also Orlandi, 2014). The visual system merely behaves *as if* it were operating under the assumption that light comes from above. If this is right, these assumptions cannot be examples of perceptual representations that function to represent in a stimulus-independent fashion.

Of course, a popular view is not necessarily a correct one. Although I am drawn to the natural constraints view, I shall not try to defend it here. Instead, I conclude with a qualification: Insofar as there are no perceptual representations that function to represent in a stimulus-independent fashion, long-term memory representations are cognitive representations. In turn, recognition is partly cognitive because it in part constitutively depends on such representations. Although our discussion has been cast in terms of Bruce and Young’s face recognition model, this conclusion generalizes to any form of recognition that in part constitutively depends on long-term memory representations (e.g., familiar object or scene recognition).

7. Conclusion

I have argued that considerations relating to perspectival sensitivity indicate that recognition (in the token sense) is at least partly perceptual. Given the constitutive role of long-term memory representations in recognition, I found reason to think that recognition is also partly cognitive. In sum, recognition is an interface capacity that straddles the border between perception and cognition.

Our conclusion bears on the views set out by Strawson, Evans, and Peacocke (Section 1). One consequence of our discussion is that a creature incapable of forming long-term memories would lack a capacity for recognition. For Evans, this creature would lack a very basic means of identifying particulars non-descriptively; for Strawson, this creature would have *no* means of identifying particulars. On Peacocke’s view, such a creature would be conceptually impoverished relative to us. It would not possess the concept “Lincoln Plaza”, or any other recognitional concept for that matter. In short, if Strawson, Evans, and Peacocke are correct, long-term memory is significant for much more than remembering—it plays a foundational role in identification and concept possession as well.

Key questions remain. How are we to understand intellectually robust forms of recognition (e.g., recognizing a familiar mathematical proof)? What are the epistemic dimensions of recognition? In the next chapter, I focus on just one question that naturally arises from the present discussion: What is the role of affect in recognition?

Chapter 2: A puzzle concerning the role of affect in recognition

1. Introduction

Recognition and feeling go hand in hand. Recognizing a favorite childhood toy in an unexpected location may be a source of surprise or joy; recognizing a bully from elementary school can stir up feelings of fear or anger. A misplaced sense of familiarity without recollection might be partially responsible for the idiosyncratic feelings that are characteristic of *déjà vu* (see Cleary, 2008). Salient feelings such as these tend to be present when something strikes us as familiar, occupying the foreground of one's experience. Other examples are more subtle, such as when one enjoys a feeling of warm comfort while driving along familiar route. One reason for this connection is that we often care about the people, places, and things we are most familiar with, so recognizing these things invariably generates an affective response.

But how exactly are we to understand the role of affect in recognition? The paragraph above hints at the idea that recognition *causes* certain affective responses. But does affect play a constitutive role in recognition? The Causalist says no. According to *Causalism*, affect is only causally implicated in recognition—either as a cause or a causal byproduct of recognition (or both).¹⁸ In contrast, the Constitutivist says yes. According to *Constitutivism*, part of what it is to recognize a person, place, or thing is to enjoy a certain affective response—a feeling inherent in the process of recognition, so to speak. To my knowledge, the distinction between Causalism and Constitutivism has not been acknowledged in the philosophical literature. In this chapter, I explore both views and highlight the difficulty in trying to assess which is correct.

To better understand the disagreement between Causalists and Constitutivists, it is useful to have an intuitive grip on the causal-constitutive distinction. A sculptor uses a chisel to shape a piece of clay. Slowly, a statue takes form. The chisel is part of the causal story as to how the statue came into existence but, unlike the clay, does not constitute it. Similarly, smoke is caused by fire but does not constitute it. Rather, it is a byproduct of the exothermic reactions between fuel and oxygen that constitute fire. Of course, not all examples are so clear cut. For instance, theorists debate whether the prefrontal cortex (PFC) partly constitutes the neural basis of consciousness or is merely causally implicated in it (Block, 2007; Michel and Morales, 2019). Those who endorse the latter view accept that the PFC partly constitutes the neural basis of verbal report, though they deny that verbal report is itself necessary for consciousness.

The disagreement between Causalists and Constitutivists bears directly on our understanding of certain clinical disorders involving recognition. I touch upon this shortly in my discussion of Capgras syndrome—a condition characterized by the delusional belief that certain familiar individuals have been replaced by imposters. It is also relevant to our understanding of artificial facial recognition systems. Consider whether your smartphone recognizes you. In simplified terms, existing facial recognition systems, such as those found in smartphones, work by matching an aggregated dataset of images of a face to a

¹⁸ Although I do not discuss this option elsewhere in the paper, the Causalist might also claim that affect is causally implicated in recognition through a common causal structure.

novel image of that same face. This is by no means an easy task, given the variation in a face's appearance over time (e.g., in lighting, facial expression). Importantly, however, feeling plays no role in this matching process. The Causalist is in a position to claim that a smartphone is capable of genuine face recognition since on their view affect is non-essential for recognition. In contrast, according to the Constitutivist, my smartphone might be able to reidentify me with some success, but it does not recognize me.¹⁹

My eventual plan is to introduce Capgras syndrome as a case study for exploring the disagreement between Causalists and Constitutivists in more concrete terms. In turn, I suggest three possible ways of resolving this disagreement and show why none are particularly satisfying. First, however, some ground-clearing is in order.

2. The causalist-constitutivist dispute

To start, what are some of the characteristics of affect, or feeling? (A clarification: I use the terms 'affect' and 'feeling' interchangeably.) For one, affective responses are often said to be *valenced*, a dimension that corresponds to their pleasantness or unpleasantness in quality. A feeling of disgust possesses a negative valence while a feeling of joy possesses a positive valence. Likewise, feelings are often said to possess a variable degree of intensity, or *arousal*. Slight annoyance, for instance, is a low intensity feeling. A number of theorists accept that all feelings possess both valence and intensity (see, *inter alia*, Reisenzein, 1994; Russell, 2003; Colibazzi et al., 2010).

Other characteristics are either less widely endorsed or fail to apply to all feelings. An example of a less widely endorsed characteristic is *dominance*. The rough idea is that feelings high in dominance, such as anger, tend to produce more controlling behavior than feelings low in dominance (so-called "submissive" feelings), such as fear.²⁰ An example of a characteristic that fails to apply to all feelings is *intentionality*, or aboutness, at least the sort of intentionality that is directed at particular people, places, and things. Kind (2014) distinguishes two types of feelings on this basis: *emotions*, which exhibit this sort of intentionality, and *moods*, which typically do not. Accordingly, undirected boredom or depression qualifies as a mood, whereas directed anger at a broken computer qualifies as an emotion.

What is recognition? A first pass answer is that it is a mental process that functions to reidentify people, places, and things on the basis of their appearance. This answer provides a useful starting point since it leaves open whether affect plays a causal or constitutive role in the mental process. Here, the relevant sense of recognition is *token*-recognition, which is to be distinguished from *type*-recognition, a process of grouping things into general categories on the basis of their appearance (see Abid, 2021). If I have never seen a particular dog before, I cannot to token-recognize it, though I can type-

¹⁹ Another possible response on behalf of the Constitutivist is to restrict the scope of their view to apply only to creatures sufficiently like us. With this restriction in place, the Constitutivist need not deny that smartphones and other artificial systems are capable of recognition. I consider a different restriction on Constitutivism in 5.3.

²⁰ Dominance plays a central role in the influential PAD model introduced by Mehrabian and Russell (1974). According to the PAD model, all feelings can be classified using the dimensions of valence (or "pleasure"), arousal, and dominance.

recognize it as a dog. Two different *types* of token-recognition are often distinguished: familiarity and recollection (Mandler, 2008; Yonelinas, 2010).²¹ Unlike recollection, *familiarity* does not involve the retrieval of contextual information about one's previous encounter with a person, place, or thing. For instance, a person might strike me as familiar even if I cannot recollect who they are, or where or when I last saw them. The debate between Causalist and Constitutivists primarily concerns the role of affect in familiarity.

The Causalist-Constitutivist dispute needs to be disentangled from a longstanding debate over whether there is a *sui generis* feeling of familiarity. The idea that there is some such feeling is found in the works of James (1890) and Russell (1921), who both speak of “the” feeling of familiarity without attempting to reduce it to some other feeling. In contrast, this idea is ridiculed throughout the later works of Wittgenstein. It is a mistake, Wittgenstein thinks, to posit some unique, elusive, and inarticulable feeling of familiarity since “there are a great many different experiences, some of them feelings, which we might call ‘experiences (feelings) of familiarity’” (1969: 181). Accordingly, *the* feeling of familiarity is a misnomer; what we in fact have is a heterogenous jumble of different feelings, such as surprise, joy, anger, etc. The dispute over whether there is a *sui generis* feeling of familiarity is orthogonal to the dispute between Causalists and Constitutivists. The former concerns the *nature* of the affective response that accompanies recognition; the latter concerns the *nature of the relation* between this affective response and recognition.

3. Phenomenological considerations

Does phenomenology support Constitutivism over Causalism? The Constitutivist might argue that the qualitative character of recognition—what it is like to recognize a familiar person, place, or thing—can only be characterized by reference to affect. What else, aside from feeling, would explain the qualitative contrast between, say, recognizing a familiar city street and seeing that same street anew? By denying that affect plays a constitutive role in recognition, the Constitutivist may claim that the Causalist fails to acknowledge that there is anything that can be properly regarded as the qualitative character of recognition. At most, the Causalist can claim recognition has a “secondhand” qualitative character that derives from the feelings causally associated with it.

The Causalist will not find these considerations compelling. For one, the Causalist may simply reject that recognition has any essential qualitative character. This accords with the earlier suggestion that artificial facial recognition systems are capable of recognition, even if we suppose there is nothing it is like to be such a system. Alternatively, the Causalist may grant that recognition has an essential qualitative character but deny that it must be characterized as a feeling. Perhaps the qualitative character of recognition consists in a perceptual experience of familiarity. We do, after all, speak of things looking familiar, sounding familiar, etc.

²¹ It is worth remarking that this taxonomy, which draws on findings from scientific psychology, deviates from our folk psychological taxonomy. For one, folk psychology does not distinguish type- and token-recognition. Moreover, within our folk psychological taxonomy, the relationship between familiarity and recognition is opaque. These are not objections to the taxonomy on offer, as long as we remain clear that it is *not* an attempt to reconstruct our folk psychological taxonomy.

This latter suggestion might be met with suspicion. Lyons (2005) denies that there is a perceptual experience of familiarity. Focusing on vision, Lyons points out that there need not be much visual similarity among the class of things that are familiar to a given subject. For example, a person on Walnut Street may be familiar to me. At the same time, Walnut Street *itself* may be familiar to me. Yet, it is clearly false that the person looks like Walnut Street. For Lyons, the upshot is that “looking familiar” is just a loose way of talking—a shorthand for referring to a feeling of familiarity.

This conclusion is premature. Lyons’s reasoning at most establishes that the experience of familiarity is unlike, say, the experience of motion or shape in that it cannot be adequately characterized at the level of *what* we perceptually experience. Lyons does not, however, eliminate the possibility that the experience of familiarity consists in a certain *way*, or manner, of perceiving things (e.g., the person on Walnut Street and Walnut Street itself).²² As an analogy, if I remove my eyeglasses, everything may look blurry. Lyons may correctly point out that there is little in the way of visual resemblance among the things I see, but this does not imply that we are somehow mistaken in supposing these things all look blurry. After all, blurry may characterize the way that these things are presented in vision. On this view, blurriness is not a feature of the world. Blurry is *how* the world is visually experienced. This is why something can look blurry, but it cannot in any literal sense look to *be* blurry (Smith, 2008). In this way, blurriness is distinguished from fuzziness. Something can look to be fuzzy, such as an image with an indistinct boundary projected on a screen.²³

The upshot is that the dispute between Causalists and Constitutivists cannot be easily resolved by appeal to phenomenology. Instead, phenomenological considerations lead to an impasse.

4. Capgras syndrome

To further illustrate the difficulty of resolving the dispute, let us consider the question of whether persons suffering from Capgras syndrome recognize their loved ones. Capgras syndrome, first described by Joseph Capgras and Jean Reboul-Lachaux in 1923, is commonly distinguished by its highly topic-specific (“monothematic” or “monosymptomatic”) delusions concerning the identity of familiar individuals.²⁴ As Ellis and colleagues put it:

The Capgras delusion is characterized by the patient insisting that others, usually those quite close emotionally, have been replaced by doubles, impostors or robots. This bizarre belief can arise as a monosymptomatic delusion held with conviction despite insight into its irrationality and unbelievability. It has been noted to follow various types of brain injury, has been reported for patients from many different

²² In representational terms, this “way of perceiving” is often referred to as a *manner of representation*, a characterization of how some representational content is represented (Chalmers, 2004). Certain naïve realist approaches also appeal to ways of perceiving (see, e.g., Campbell, 2016).

²³ For an overview of various accounts of blurry vision, see Allen (2013).

²⁴ Although most Capgras delusions concern the identity of familiar persons, this is not always the case. For instance, some Capgras delusions concern familiar animals (Darby and Caplan, 2016).

cultures, and presents a significant risk of violence against the alleged impostors (1997: 1085, citations omitted).

The most widely accepted explanation of Capgras syndrome—at least instances that are confined to the sense modality of vision—is that the disorder (in part) originates from a failure to form appropriate affective responses to the faces of familiar individuals (Ellis and Young, 1990; Ellis et al., 1997; Hirstein and Ramachandran, 1997).²⁵ The strongest evidence for this claim is that the skin conductance responses (SCRs) of visual Capgras patients to the faces of familiar and unfamiliar individuals—a measurement commonly used to assess affective responses to stimuli—are virtually indistinguishable. These same patients show normal SCRs to familiar and unfamiliar auditory tones, so it cannot be that these patients simply show no differential SCRs to familiar and unfamiliar stimuli in general. Rare instances of Capgras syndrome that are isolated to the auditory modality are explained in terms of a failure to form appropriate affective responses to the voices of familiar individuals (Lewis et al., 2001). Moreover, in patients with auditory Capgras syndrome, SCRs to the voices—but not faces—of familiar and unfamiliar individuals are virtually indistinguishable.

Assuming this widely accepted explanation is correct, Causalists will be able to claim that visual Capgras patients *can* recognize the faces of familiar individuals through sight alone despite failing to form an appropriate affective response. This would imply that the disrupted response of the Capgras patient is post-recognitional. In support of this claim, Capgras patients tend to perform moderately well on facial reidentification tasks (Hirstein and Ramachandran, 1997; Ellis et al., 1997), such as the Warrington Recognition Memory Test in which patients are simultaneously shown images of two distinct faces—one novel and one displayed along with dozens of others at an earlier point—and are forced to decide which they have seen before.

Constitutivists will insist that Capgras patients *cannot* recognize the faces of familiar individuals through sight alone since, on their view, forming an appropriate affective response is part of what it is to recognize a familiar face. Intuition might seem to favor Constitutivism. It is odd to claim that Capgras patients recognize their loved ones but regard them as imposters. Yet, this view has the consequence that moderately successful performance on a facial reidentification task is possible even if one is incapable of recognizing a familiar face.²⁶ Who is correct?

5. Resolutions

I can imagine three possible resolutions to this dispute. First, it may be that the new or existing evidence puts the matter to rest. Second, it may be revealed that the dispute is

²⁵ According to “two-deficit” accounts of delusion formation (Langdon and Coltheart, 2000; Coltheart, 2010), visual Capgras delusions arise from an impaired affective response to faces along with an impaired capacity for doxastic evaluation.

²⁶ In defense of Constitutivism, this consequence is not as outlandish as it initially seems. For instance, even if Capgras patients are incapable of recognizing familiar faces, they may still complete the Warrington Recognition Memory Test by reidentifying idiosyncratic facial features (e.g., short noses, pronounced ears) rather than by reidentifying faces as such. As I discuss shortly, the Causalist can accommodate the possibility that Capgras patients are incapable of recognizing familiar faces.

merely verbal. Finally, and related to the previous option, it may be that the dispute is settled by disambiguating the relevant sense of recognition. In this section, I discuss each of these possible resolutions in turn and show why none are particularly promising.

5.1 New or existing evidence

It may be that new or existing evidence concerning the role of affect in recognition settles the matter. I am doubtful that the dispute will be resolved in this way since any such evidence will presumably be open to two mutually inconsistent interpretations. Which interpretation is regarded as the more plausible one will depend on one's antecedent position in the debate. For example, suppose that new neuroimaging evidence indicates that activity in neural region X (e.g., a part of the limbic system) correlates with the affective responses associated with recognition. The Constitutivist will presumably interpret this finding as an indication that region X belongs to one of the neuroanatomical pathways for recognition (see 5.3 below), while the Causalist will instead suggest that region X lies outside any such pathway. Anatomical distinctions reflect one's theoretical commitments.

Similar considerations apply in the case of behavioral evidence. Suppose it is shown that subjects have an improved ability to recognize friends in low-resolution photographs when they are in a good mood.²⁷ At first blush, such a result would seem to support Constitutivism. Yet, the Causalist could offer their own interpretation of the finding according to which a subject's affective state is dissociable from the process of recognition but nonetheless exerts a causal influence on it—what amounts to a kind of top-down feedback.

Indeed, the Causalist might extend this interpretation to provide a more nuanced explanation of Capgras syndrome than the one given in §4. As I mentioned earlier, Capgras patients tend to perform moderately well on facial reidentification tasks. However, this is not to say that they perform *equally* well as control subjects; in fact, at least one study indicates they do not (Fiacconi et al., 2014; see also Lucchelli and Spinnler, 2008). In light of this disparity in task performance, the Causalist may grant that Capgras patients suffer from an impairment in familiar face recognition but claim that an absence of affective feedback explains why this is so.²⁸

Just as both the Causalist and Constitutivist can provide competing interpretations of the influence of affect on recognition, both can offer competing interpretations of the influence of recognition on affect. Consider the *mere exposure effect*: Familiar stimuli are usually evaluated more positively than novel stimuli (Zajonc, 1968). For instance, we tend to find a face more attractive if we have seen it before. While it may seem that the mere

²⁷ This hypothetical result would differ from existing evidence indicating that subjects are more likely to *mistakenly* judge that a stimulus is familiar when they are in a good mood (Garcia-Marques et al., 2004). An explanation of this latter finding is that subjects tend to use positive affect as a heuristic guide to familiarity, especially in situations of uncertainty (Monin, 2003).

²⁸ An explanation along these lines is suggested by Fiacconi et al. (2014). This disparity in task performance may also be explained by a perceptual or memory deficit unrelated to affect. For evidence for the latter possibility, see Darby and Sashank (2016).

exposure effect is grist to the mill of the Constitutivist, the Causalist will contend that the phenomenon can be understood in causal terms: positive affect is a causal byproduct of recognition. Hence, a competing interpretation is again available to the Causalist.

5.2 A verbal dispute

A second possibility is that it may become apparent that the dispute is a verbal disagreement concerning the correct use of the terms ‘recognize,’ ‘recognition,’ etc. Though I find this resolution more plausible than the last, it certainly seems as though the issue is non-verbal—that there is an objective fact of the matter, independent of the way in which we use our words, as to whether Capgras patients recognize their loved ones. In any case, even if the dispute partly concerns our use of words, it does not *wholly* concern our use of words. The dispute has implications for the kinds of interpersonal relationships Capgras patients are able to enter into, whether they are responsible for acts of violence towards alleged imposters, the evidential status of a witness testimony they might provide in a court of law, their knowledge of the past, and so on. These implications are arguably all non-verbal. In contrast, wholly verbal disputes typically lack non-verbal implications.

As an aside, even if we *were* to conclude that the dispute is wholly verbal, this would not in itself guarantee a resolution to the debate by revealing it to be non-substantive. While many wholly verbal disputes are non-substantive, it is doubtful that all are. In certain fields, such as linguistics or philosophy of language, substantive disputes are almost entirely over words (Chalmers, 2011). Examples can be found within the philosophy of mind as well, such as disputes over how we ought to understand the term ‘representation.’ In my mind, these disputes are verbal, but they are also substantive.²⁹ In opting for a particular interpretation of the term, theorists are attempting to track an explanatorily significant notion within cognitive science—they are not merely expressing arbitrary preferences. We cannot rule out the possibility that a parallel point applies when it comes to the terms ‘recognize,’ ‘recognition,’ etc. The general lesson is that it is a mistake to assume a dispute will evaporate just because it concerns language.

5.3 Reconciliation through disambiguation

Related to the last option, reconciliation might be achieved through disambiguation. The most promising suggestion along these lines comes from Ellis and Young (1990), who distinguish between conscious “overt” recognition associated with a ventral stream pathway, which runs from the brain’s occipital to temporal lobes, and unconscious “covert” recognition associated with a dorsal stream pathway, which runs across the occipital to parietal lobes. They hypothesize that overt recognition is preserved in Capgras syndrome, while covert recognition is impaired because of damage to the dorsal stream (see also Ellis and Lewis, 2001). Importantly, they claim that processing within the dorsal stream is what imbues one’s impression of a face with an “affective tone” that is missing in

²⁹ To be clear, not all theorists characterize these disputes as verbal (see, e.g., Orlandi, 2014, pp. 117-120). It seems to me that part of the hesitancy in regarding these disputes as verbal stems from conflating verbal with non-substantive disputes.

the case of Capgras syndrome. This hints at the possibility that Constitutivism may be correct if restricted to just covert recognition.

To be clear, in claiming that the Capgras patients cannot imbue faces with an affective tone, Ellis and Young are *not* committed to claiming that Capgras patients experience no affective response whatsoever as a result of seeing their loved one. In Capgras syndrome, there is a hypothesized mismatch between the outputs of the overt recognition system (“I recognize them”) and the covert recognition system (“I do not recognize them”). It is possible that this mismatch induces an anomalous affective response in the patient upon seeing their loved one, such as a feeling of *unfamiliarity*.

In any case, there are difficulties with the suggestion that Constitutivism may be correct if restricted to just covert recognition associated with a dorsal stream pathway. First, it presupposes that information processing in the ventral stream is responsible for the character of one’s conscious experience while information processing in the dorsal stream is not. This is a widely held presupposition, but some deny it. Wu (2014), for instance, argues that information processing in the dorsal stream is responsible for the experience of a stable world despite quick, pervasive movements of our eyes that result in constant changes in the retinal image—a phenomenon known as “space constancy.” (To appreciate the *failure* of space constancy, gently press on your eye with a finger and notice how everything seems to move.) If this is right, it is a mistake to regard the dorsal stream as an unconscious “zombie” stream.

Bracketing this concern, the Constitutivist may worry that this reconciliation is too concessive. By restricting Constitutivism to just unconscious covert recognition, we undercut the phenomenological considerations that the Constitutivist raised against the Causalist in §3, namely, that affect must play a constitutive role in recognition if we are to explain the qualitative character of recognition. This point, after all, pertains only to conscious overt recognition.

These first two difficulties might be brushed aside. In response to the first difficulty, even if we are skeptical of the idea that the dorsal stream is responsible for unconscious processing while the ventral stream is responsible for conscious processing, we can still make sense of a non-anatomical distinction between a covert recognition system associated with affective responses and an overt recognition system that is not, which is all that is needed to articulate the restricted version of Constitutivism. In response to the second difficulty, we have already seen that phenomenological considerations are not effective against the Causalist, so no argumentative ground would be lost if the Constitutivist were to abandon them.

Yet, there is a deeper difficulty that cannot easily be surmounted: One worries that the disagreement between Causalists and Constitutivists will simply resurface after we restrict our attention to covert recognition. That is, the role of affect in covert recognition will itself be a source of debate. The Causalist will claim that affect is only causally implicated in covert recognition, whereas the Constitutivist will claim that it plays a constitutive role. No argumentative advances are made by the restriction to covert recognition. In fact, casting the disagreement in these terms introduces new epistemic difficulties. Unconscious perception is notoriously difficult to study; theorists do not even

agree whether it exists (Peters et al., 2017). If unconscious perception is some indication, we should expect serious hurdles when it comes to the study of unconscious covert recognition. For this reason, reconciliation through disambiguation is ultimately a non-starter.

6. Conclusion

While the disagreement between Causalists and Constitutivists is significant and arguably non-verbal, as our discussion of Capgras syndrome helps illustrate, it is far from obvious how we should go about resolving it. Neither phenomenology nor new empirical evidence appears capable of advancing the debate beyond a stalemate. Likewise, the most obvious route to reconciliation through disambiguation looks to be a non-starter. We are therefore left with a puzzle concerning the role of affect in recognition.

Analogous puzzles seem to arise in other domains involving affect. For example, is the feeling of knowing causally or constitutively implicated in knowledge? Is a feeling of effort causally or constitutively implicated in action? Should the relationship between affect and perception be understood in causal or constitutive terms? It remains to be seen just how closely these cases align with the present discussion. Still, if what I have said is any guide, providing answers to these questions will be no straightforward task.

In this and the previous chapter, I have investigated whether perception, cognition, and affect play a constitutive role in recognition. However, I have not yet addressed the question of what exactly it *is* to recognize a particular person, place, or thing. This—and closely related questions concerning the explicability and analyzability of recognition—are the focus of the next chapter.

Chapter 3: The irreducibility of recognition

1. Introduction

In this chapter, I focus on the question of what it is for a subject to recognize a particular. Some examples of this phenomenon are vivid and infrequent, such as when you recognize a lake you have not visited since you were a child. Other examples are mundane and pervasive, such as when you recognize your scarf hanging on the restaurant coat rack. Also notable are our failures to recognize, for instance, when your colleague in the hallway fails to recognize you and walks by without saying a word. While these examples all involve sight, other sense modalities also do the job (e.g., when we recognize a loved one by the sound of their voice or a favorite trinket by touch). Here, the relevant sense of recognition is one that involves the reidentification of particulars, such as people, places, and things. This sense of recognition is to be distinguished from the recognition of people, places, and things *as such*, for example, when you recognize the creature in the bushes *as a tiger*.³⁰

Aside from its intrinsic interest, the question of what it is to recognize a particular is worth exploring because it appears to be a precondition for other important aspects of thought. Thus, it has been argued that the capacity for recognition is necessary for the identification of particulars (Strawson, 1959; cf. Evans, 1982), non-probabilistic forms of reasoning (Evans, *ibid.*), a conception of the *same* material object extended over time (Swinburne, 1969), and the possession of a range of concepts, including locational concepts, such as the concept *Lincoln Square*, and demonstrative concepts, such as the concept *that shade* (Peacocke, 1992 and McDowell, 1994, respectively). A common thread underlying these arguments is that our capacity for recognition plays a foundational role in our mental economy. In due course, I argue for a more precise version of this idea.

The question is also connected to the work of cognitive scientists who have explored the subpersonal underpinnings of recognition. Perhaps the most influential contribution along these lines is Bruce and Young's (1986) familiar face recognition model. On this box-and-arrow model, the face of an individual is inputted into a system that produces a *structural code* of that face, in essence, a description of that face stripped of its facial expressions. This newly produced structural code is matched against old codes stored in dedicated systems called *face recognition units*. The rough idea is that if a match is made, the face is treated as familiar by the system, though this can be overridden by a decisional mechanism if the face is judged to belong to, say, a friend's never before seen doppelgänger.³¹ As we will see, there is a pervasive inclination to extend such "matching" explanations to the personal level. A secondary aim of this chapter is to diagnose why such explanations are destined to fail.

I begin by considering an analysis of recognition in terms of a relation of re-acquaintance (Section 2). Even with supplementation, this analysis cannot explain mundane cases in which

³⁰ Are these two different senses of recognition dissociable from one another? The answer seems to be yes. You might recognize the tiger in the bushes as a tiger without recognizing it in the reidentificational sense (e.g., if you have never encountered that specific tiger before). Conversely, you might recognize the tiger in the bushes in the reidentificational sense even if you do not recognize it as a tiger (e.g., you might think it is a bear). For a more detailed discussion of the dissociation between these two different senses of recognition, see Abid (2021).

³¹ For a discussion of recognition in light of the possibility of unfamiliar duplicates, see Evans (1982, pp. 278-284).

standing in the re-acquaintance relation with a particular does not suffice for recognizing it. What is missing from this first analysis is an internal connection between one's past and present encounters with a particular. I turn to an alternative analysis which seeks to explain recognition using the framework of mental files (Section 3). While this second analysis avoids the problems faced by the first, it nonetheless fails to explain instances of recognition that involve familiarity without recollection.

The idea that recognition eludes explanation becomes the dominant theme in the remainder of the chapter. Inspired by the remarks of Wittgenstein and Evans, I introduce what I call the *Inexplicability Thesis*, according to which recognition cannot be mentalistically explained (Section 4). After unpacking the relevant notion of mentalistic explanation, I argue that the Inexplicability Thesis leaves us with a meta-explanatory problem: We need an explanation for why we cannot provide a mentalistic explanation of recognition (Section 5).

I develop an ontological solution to this meta-explanatory problem according to which recognition cannot be subjected to mentalistic explanation because it is an irreducible mental capacity. My solution respects the observation that our capacity for recognition depends on more basic perceptual capacities for detection and discrimination (Section 6). Nevertheless, if I am right, there is no X such that the capacity for recognition *fully* depends on the capacities for detection, discrimination, and X (where X need not be just one mental capacity). I show how my solution can be established non-circularly by treating the irreducibility of recognition as an empirical hypothesis (Section 7). I conclude by briefly discussing *non*-mentalistic explanations of recognition (Section 8).

2. Recognition as re-acquaintance

It might initially seem that recognition is simply a matter of observing a particular twice over. To see why this suggestion is overly simplistic, it is instructive to consider the following first-pass proposal:

Recognition as Re-Acquaintance: A subject S recognizes some particular o just in case S is re-acquainted with o .

Let us say that S is *re-acquainted* with o just in case S was perceptually acquainted with o in the past and is perceptually acquainted with o again in the present.³² In this context, we should adopt a liberal conception of perceptual acquaintance to allow for cases of indirect perception in which, for instance, S recognizes a celebrity on the street despite having only previously seen them on television.

As it stands, there are obvious counterexamples to Recognition as Re-Acquaintance. Suppose that S 's memory of o is totally wiped because of a severe accident. With no memory of o whatsoever, S will not recognize o . Yet, this complete eradication of memory concerning o will not prevent S from standing in the re-acquaintance relation with o .

³² For a discussion of the relation of re-acquaintance, see Recanati (2012, pp. 86-88). While some of Recanati's remarks suggest that he endorses Recognition as Re-Acquaintance, I set these exegetical issues aside for now and focus on discussing the proposal in its own right.

Counterexamples persist even if we restrict ourselves to ‘well-functioning’ subjects—whatever exactly that might amount to. Suppose that some well-functioning subject *S* stands in the re-acquaintance relation with some particular *o*. Suppose further that *o* has changed some of its superficial appearance properties—properties that are inessential to sustaining *o*’s numerical identity over time. For instance, if *o* is a person, we might imagine that *o* has changed their hair color, grown facial hair, developed a few wrinkles, etc. (On *any* non-eliminativist theory of personal identity, *o* can retain their numerical identity despite such changes.) If this change in appearance properties is drastic enough, *S* will no longer be able to recognize *o*. To take an extreme example, if *o* is wearing thick face paint, *S* will not be able to recognize *o* despite standing in the re-acquaintance relation with them.

In fact, it is not even essential that *o* change its superficial appearance properties. Suppose that *S* is an untalented astronomer and that *o* is a star clearly visible in the night’s sky. Even after a great deal of practice, *S* may be unable to recognize *o*. (“All the stars look the same to me!”) Yet, all the conditions necessary for *S* to stand in the re-acquaintance relation with *o* might be met: *S* may have clearly perceived *o* in the past, guaranteeing prior perceptual acquaintance with it, and stand in the very same acquaintance relation with *o* in the present.

The problem, in its most general form, is that a subject *S* can stand in the re-acquaintance relation to a particular *o* without reidentifying *o*. However, *S* cannot recognize *o* *unless* *S* reidentifies *o*. Hence, *S* may stand in the re-acquaintance relation to *o* but fail to recognize *o*.

Putting things this way suggests that the problem might be avoided if we add a reidentification condition to the re-acquaintance relation. The rough idea would be that a subject counts as *reidentifying* a particular if they in some way register that they have observed it before. Let us say that a subject *S* stands in the *re-acquaintance** relation to a particular *o* just in case *S* stands in the re-acquaintance relation with *o* and, in addition, reidentifies *o*.

This is not an adequate fix. Standing in the re-acquaintance* relation is still not sufficient for recognition. This is because reidentification is itself not sufficient for recognition (although it is necessary). To see this, note that the untalented astronomer may reidentify the North Star even if they are unable to recognize it; they might, for instance, reidentify it by asking a professional astronomer to point it out or by using a mobile phone app. Similar considerations apply in the case of a friend wearing face paint that makes them unrecognizable. Recognizing something is just one way of reidentifying it.

In these cases of reidentification without recognition, the issue is that there is no internal connection between *S*’s past and present acquaintance with *o*. In such cases, *S* perceives *o* and has a memory of *o*, but fails to link their perception and memory of *o*. This raises the question: What exactly does it mean to ‘link’ one’s percept of *o* with their memory of *o*? To provide an answer to this question, I turn to an attempt to analyze recognition using the framework of mental files.³³

³³ A different approach takes inspiration from the causal theory of seeing (Grice, 1961). According to the simple version of the theory, *S* sees *x* just in case *x* causes *S* to have a visual experience as of *x*. Perhaps an analogous theory could be constructed for recognition by building on the notion of re-acquaintance: *S* recognizes *x* just in case *x* causes *S* to have a recognitional experience as of *x* when *S* is re-acquainted with *x*. Now, there are well-known counterexamples to the simple causal theory of seeing involving deviant causal chains. For example, suppose *S* is in

3. Recognition via linking

Many theorists have been drawn to the idea that a file-like structure pervades our minds. So-called ‘mental files’ have been invoked to provide psychologically realistic accounts of continued belief over time, informative identity statements, etc. This raises the question: Could mental files be used to analyze recognition?

Consider the following approach:

Recognition via Linking: A subject *S* recognizes some particular *o* just in case a perceptual file of *o* initiated by *S* at time *t* is linked to a memory file of *o* initiated by *S* at some earlier time *t*'.³⁴

Perceptual and memory files are both types of *mental files*. While much has been said about mental files, the important point for our purposes is that a mental file about *o* contains—but is itself distinguishable from—a cluster of information taken to correspond to *o* and *o* alone.³⁵ We might think of a *cluster of information* about *o* as a set of predicates all taken to apply to *o* or, slightly more figuratively, as a body of documents all taken to pertain to *o*.³⁶ Regardless, once a mental file is created, it comes to function as a repository for information specifically about *o*. If the cluster of information contained in a mental file about *o* derives from *S*'s present perceptual acquaintance with *o*, the mental file is a *perceptual* file. If the cluster of information instead derives from *S*'s past acquaintance with *o*, the mental file is a *memory* file.

When two files are *linked*, information flows freely from one file to the next. Put differently, during the period in which two files are linked, it is *as if* their information is contained in a single file. Linking is a transient operation. Paradigmatic examples of linking involve the acceptance of an informative identity. If I accept that Clark Kent is Superman, my CLARK KENT and SUPERMAN files get linked. Consequently, informational integration across

the immediate presence of *x* but *S*'s visual experience as of *x* is a result of the manipulations of a neurosurgeon, as opposed to being caused by *x* as would normally be the case (Noë, 2003). And suppose that the neurosurgeon does this *because* of *x*. For example, a manipulative neurosurgeon might see *x* in the immediate presence of *S* and subsequently devise a plan to create a visual experience in *S* indistinguishable from one normally caused by *x*, thereby leading to a causal chain from *x* to *S* mediated by the neurosurgeon. Intuitively, *S* does not see *x*, but *S*'s visual experience as of *x* is nonetheless caused by *x* (through a deviant causal chain). Counterexamples like these, which require an articulation of the ‘right sort’ of causal chain, plague the analogous causal theory of recognition. An additional problem specific to the causal theory of recognition is that it is unclear what a “recognitional experience” might amount to. The most promising candidate is a feeling of familiarity. This feeling does accompany some instances of recognition. For instance, I enjoy a feeling of familiarity when I recognize a colleague by surprise in a foreign city. However, some hypothesize that the feeling of familiarity occurs only if there is a violation of expectation (Whittlesea and Williams, 2000). This hypothesis seems to correctly predict that a feeling of familiarity would not be stirred up if I were to recognize a colleague in the hallway at work—just where I expect to encounter them. Similarly, it would seem to vindicate Wittgenstein’s (1958/1969) skepticism that a feeling of familiarity *must* be stirred up every time someone, for instance, gazes at their old coat and hat (a context in which there is no violation of expectation). If this hypothesis is correct, it would be misguided to think that a recognitional experience consisting in a feeling of familiarity accompanies each instance of recognition.

³⁴ A predecessor of this approach can be found in the works of Strawson (1974, pp. 46-7).

³⁵ This claim is accepted even by theorists who deny that the referent of a file is fixed by the descriptive information contained in that file (see Perry, 1980 and Recanati, 2012).

³⁶ See Perry (1980) and Schroeter (2007), respectively.

files becomes possible. Suppose that the predicate IS A SUPERHERO is contained in my SUPERMAN file and that the predicate LOVES LOIS LANE is contained in my CLARK KENT file. If these two files are linked, I will be able to infer that a superhero loves Lois Lane. Note that my CLARK KENT and SUPERMAN files do not collapse into a single file after being linked.³⁷

There are a few reasons for having doubts about Recognition via Linking. Some argue that it conflicts with a plausible criterion for token-individuating mental files.³⁸ My concern is that it cannot explain cases of recognition that involve familiarity without recollection.

Familiarity is often contrasted with *recollection*, which requires the retrieval of contextual information about a particular. In contrast to recollection, a particular can strike one as *familiar* even if one cannot pinpoint when or where one last observed it. Consider the classic butcher in the bus example (Mandler, 1980). I get on a bus and notice a seemingly familiar person sitting in the back. Even though I am sure I have seen this person before, I cannot recollect who this person is, where I last saw them, etc. It is only *after* the individual leaves the bus that I can recall who they are—my butcher.

How might we make sense of that moment in which I could not recollect anything about the person on the bus even though they struck me as familiar? It cannot be in terms of the linking of perceptual and memory files. If these files *were* linked, I would be able to freely integrate information contained in my memory file with information contained in my perceptual file. I would, for instance, be able to infer that the person sitting on the back of the bus cuts meat. Yet, this is precisely what does not happen in the example above. I cannot recall even the most rudimentary information corresponding to my previous encounter with the person.³⁹

Nevertheless, such information is clearly stored in my memory file of the individual. A simple prompt might be all I need to access a great deal of information about them. For instance, if you were to remind me that the individual is my butcher, I might reply, “Yes, of course, now I remember; they were wearing a blood covered smock the last time I saw them.” Therefore, it cannot simply be that the cluster of information contained in my memory file is empty.

Could the problem be one of information retrieval? Perhaps my perceptual and memory files *are* linked even though I am unable to retrieve information about my past encounters with the individual. It is hard to make sense of this suggestion. Again, if my perceptual and memory files *were* linked, there would be free information flow between the two files. Consequently, we would expect a *symmetry* in the retrieval of information from each file: I would either be (i) able to retrieve information from both files or (ii) unable to retrieve information from either file. What we in fact find is an *asymmetry* in retrieval. On the one hand, I can retrieve some information about the individual, namely, information contained in my perceptual file (e.g., that the person sitting in the back of the bus is wearing blue). So, (ii) is not an option. On the other hand, I cannot retrieve information about my past encounters with the individual, information housed in

³⁷ See Lawlor (2001); Recanati (2012).

³⁸ See Recanati (2012, pp. 84-5).

³⁹ It is for this reason that there is little use in appealing to what we might call *proto-linking*, an operation between two distinct files allowing for a *restricted* flow of information from one to the other. What is distinctive about cases of familiarity without recollection is that there is *no* information flow between files.

my memory file. So, (i) is also not an option. Against the present suggestion, the most plausible reason why this asymmetry exists is that my perceptual and memory files are *not* linked.

Could it be that familiarity is an altogether different phenomenon than recognition? If so, cases of familiarity without recollection would not be a threat to Recognition via Linking. Although it might seem that this suggestion is vindicated by an appeal to ordinary language, this is incorrect. While we sometimes speak of familiarity and recognition contrastively (e.g., “I didn’t recognize them, but they still seemed familiar”), this is not always the case (e.g., “I recognized a familiar face in the crowd”). What should we make of this? A natural interpretation is that the term ‘recognize’ (in its colloquial sense) is polysemous. In the contrastive cases, we use the term restrictively, requiring a subject to recollect their past encounter with a particular to count as having recognized it. In the non-contrastive cases, we use the term more permissively, allowing for familiarity-based recognition (without recollection). But if this is right, then appeals to ordinary language cannot delineate the underlying phenomena.

A more promising route would be to appeal to existing psychological models as a guide. Here, we find that matters are more conclusive: The suggestion that familiarity is a distinct phenomenon from recognition straightforwardly conflicts with widely accepted “dual-process” models of recognition (Mandler, 2008; Yonelinas, 2010). These models understand recollection and familiarity as two processes that can each give rise to recognition: You can recognize a person, place, or thing by recollecting your past encounter with it (recollection-based recognition), or you can simply recognize a person, place, or thing as familiar (familiarity-based recognition). If these models are correct, then cases of familiarity without recollection are genuine instances of recognition and hence counterexamples to Recognition via Linking.

In sum, Recognition via Linking provides only a partial explanation of recognition.⁴⁰ This is no accident: There is a sense in which recognition can never be fully explained. To bring this point into focus, it will help to turn to the work of Wittgenstein and Evans.

4. The inexplicability thesis

For both Wittgenstein and Evans, a tempting yet ultimately false explanation of recognition involves a process of matching one’s past and present impressions of a particular—a process akin to comparing an object with its picture by holding them side-by-side. Wittgenstein writes:

It is easy to misconceive what is called “recognizing”; as if recognizing always consisted in comparing two impressions with one another. It is as if I carried a picture of an object with me and used it to identify an object as the one represented by the picture. Our memory seems to us to be the agent of such a comparison, by preserving a picture of what has been seen before, or by allowing us to look into the past (as if down a spyglass) (1953/2009, §604).

⁴⁰ Perhaps in that moment of familiarity without recollection I entertain a thought of the form: *There exists a memory file of that person sitting in the back of the bus*. While this suggestion does not vindicate Recognition via Linking, it does hint that Mental File Theory might be able to account for cases of familiarity without recollection. The challenge for this approach is to specify the conditions under which I would entertain this existential thought. Presumably, I would need to stand in some appropriate relation to the person in question. What would this relation consist in? The discussion in Section 2 makes clear that spelling out this relation it is non-trivial.

Let us call this process of comparing past and present impressions *impression matching*. By invoking impression matching to explain recognition, we arrive at:

Impression Matching: A subject *S* recognizes a particular *o* just in case *S* matches past and present impressions of *o*.

Both Wittgenstein and Evans find Impression Matching implausible. For Wittgenstein, Impression Matching is phenomenologically inadequate. In recognizing a particular, it is not as though one has a sense of a side-by-side comparison of two distinct impressions (ibid., §605). For Evans, problems arise because Impression Matching fails to explain cases in which a particular is familiar to us even though we cannot recollect our last encounter with it, such as in the butcher in the bus example (Section 3). It is difficult to see how Impression Matching is supposed to apply in these cases. For in such circumstances a subject lacks access to their past impressions of the particular in question (e.g., the butcher), precluding a comparison of past and present impressions. Evans's worry is not so different from the worry raised against Recognition via Linking, suggesting that Recognition via Linking is a descendant of Impression Matching.

Although it seems to me that these criticisms are largely successful, I do not consider them in detail. My primary interest is instead in the alternative that emerges in light of these criticisms. Inspired by Wittgenstein's remarks, Evans writes:

Only the conviction that it [Impression Matching] must be so, that the process of recognition would be incomprehensible unless it were so, could blind one to the fact that it is not so. Perhaps, in some sense or other, information about the object's appearance is stored in the nervous system, but this is not information which the *subject* has, or in any sense *uses* to effect an identification [viz., a reidentification]... *We need not use anything to make an identification... the philosophical resistance to the idea that there are certain things we just do, for which a neural, but not mental, explanation should be provided, are brilliantly analysed by Wittgenstein...* (1982, p. 288, italics added).

Evans does not explicitly claim that recognition *cannot* be explained in this special mentalistic sense, only that it *should* not. I find it quite plausible that the reason why Evans thinks we should not offer such an explanation is that we cannot—it is a forlorn cause. Regardless of whether this interpretation is correct, the important point for our purposes is that Evans's remarks suggest a thesis that is independently interesting:

The Inexplicability Thesis: Recognition cannot be mentalistically explained.

Before evaluating the Inexplicability Thesis, let us take a moment to clarify the relevant sense of mentalistic explanation. Distinguishing mentalistic explanation from neighboring explanations will allow us to acknowledge that there are non-mentalistic forms of explanation applicable to recognition. I turn to the issue of what would justify the Inexplicability Thesis in later sections.

To begin, consider the finding that cortical activity in the parahippocampal place area (PPA) and the retrosplenial cortex (RSC) is closely correlated with scene recognition (Epstein and Higgins, 2007). We cannot dismiss these regions as explanatorily irrelevant to scene recognition, given that they may ultimately play a key role in an explanation of its neural basis.

This is to say that denying that there can be a mentalistic explanation of recognition must be compatible with accepting that there can be *neural* explanation of it, a point that Evans explicitly grants in the passage above. The difference is that while mental phenomena are a part of the explanandum of both a neural and mentalistic explanation, they are *only* part of the explanans of a mentalistic explanation.

Denying that there can be a mentalistic explanation of recognition must also be compatible with accepting that there can be a subpersonal explanation of recognition cast in information-processing terms—the sort of explanation a cognitive psychologist might provide (cf. Section 1). Evans implicitly leaves room for such an explanation, admitting that “If we were programming a computer to recognize faces, we should programme it to be sensitive to the ratios of the distances between various facial landmarks...” (1982, p. 290). Nevertheless, information about facial landmark ratios would be quite alien to a subject *qua* subject. This much is evidenced by the fact that a subject would be surprised to learn that they possessed such information, even if they were able to comprehend it. Unlike a cognitive psychological explanation, the explanans of a mentalistic explanation is couched at the personal level.⁴¹

Finally, denying that there can be a mentalistic explanation of recognition must be compatible with accepting that we can specify its causal role with respect to other mental phenomena. If Lewis (1972) is correct, we can always analytically define a mental state (event, process, etc.) as the occupant of a unique causal role provided by folk psychology (viz., as the x such that x causes such-and-such behavioral responses, is causally related to such-and-such other mental states, and is caused by such-and-such sensory stimuli). If such functional definitions are counted as mentalistic explanations, we would render the claim that recognition cannot be subjected to mentalistic explanation trivially false. The way out of this problem is to note that a mentalistic explanation, unlike a functional definition, is a species of non-causal explanation.

Putting these threads together, a *mentalistic explanation* is a non-causal explanation that takes a mental phenomenon as its explanandum and includes other mental phenomena as a part of its personal-level explanans. When Locke seeks to explain a complex idea in terms of a combination of simple ideas, he is providing a mentalistic explanation. Impression Matching and Recognition via Linking are both *attempts* to provide mentalistic explanations of recognition.⁴²

5. The meta-explanatory problem

⁴¹ By invoking the notion of a mental file, one might worry that Recognition via Linking falls short of this requirement for mentalistic explanation. I think this worry is misplaced. Most theorists hold that mental files play a role in our cognitive economy at the personal level, though they occasionally appeal to the subpersonal level to explain how mental files work (see, *inter alia*, Jeshion, 2010; Recanati, 2012, ch. 8; Papineau, 2013; Goodman, 2016). At the subpersonal level, the relevant notion is that of an *object* file. For a recent discussion of object files, see Green and Quilty-Dunn (2021).

⁴² Additional examples of mentalistic explanation can be found throughout contemporary philosophy of mind. Consider higher-order theories of consciousness. As Gennaro writes, “[higher-order] theories attempt to explain consciousness in *mentalistic* terms, that is, by reference to such notions as ‘thoughts’ and ‘awareness’” (2004, p. 2). On standard versions of the higher-order theory, consciousness is mentalistically explained in terms of the interaction between two nonconscious mental states: a mental representation R and a higher-order representation of R . As a second example, consider theories which attempt to explain concepts in terms of their possession conditions (see, e.g., Peacocke, 1992). Such theories can be understood as providing a mentalistic explanation of individual concepts in terms of the capacities of a thinker who has mastered them.

While the failure of Impression Matching and Recognition via Linking might *motivate* us to accept the Inexplicability Thesis, a rejection of the former does not *justify* an acceptance of the latter. Without any supplementary reasoning, it would be fallacious to infer that there could be no mentalistic explanation of recognition from the isolated failure of these two attempts. After all, what assurance do we have that there is no other mentalistic explanation of recognition—one that avoids the objections raised thus far?

To put matters slightly differently, by pointing out the problems with Impression Matching and Recognition via Linking, we provide negative arguments in favor of rejecting specific mentalistic explanations of recognition. However, these arguments do not speak to the independent plausibility of the Inexplicability Thesis. This puts the Inexplicability Thesis in a dialectically weak position, making its adequacy dependent on the inadequacy of every other possible mentalistic explanation of recognition.

Might we instead opt for the non-modal claim that recognition *has not yet been* mentalistically explained—in contrast to the modal claim that recognition *cannot be* mentalistically explained? The problem is that this non-modal claim looks much more like a sociological observation than a philosophical position, and it is vulnerable to the possibility of forthcoming mentalistic explanations of recognition.

If the Inexplicability Thesis is to be retained, we need an explanation (or better: a meta-explanation) for why there could be no mentalistic explanation of recognition. Call this the *meta-explanatory problem*.

6. The recognitional capacity as bedrock

Broadly speaking, there are two strategies for solving the meta-explanatory problem. One strategy would be to identify some epistemic basis for why we cannot mentalistically explain recognition. Maybe some idiosyncratic feature of the human mind makes it the case that we cannot provide a mentalistic explanation of recognition, leaving us cognitively closed off with respect to it.⁴³ Such an approach seeks to provide an *epistemic* answer to the meta-explanatory problem. A different strategy—the one I shall pursue—would be to demonstrate that no mentalistic explanation of recognition can be provided because recognition is, in some sense, mentally foundational. This latter strategy aims to provide an *ontological* answer to the meta-explanatory problem.

What would it mean to be mentally foundational? We can provide an answer to this question by drawing on the notion of a mental capacity. A *mental capacity* is a kind of mental tool which possesses some function.⁴⁴ We can distinguish the possession of a mental capacity from its exercise—the latter being a mental activity or process. A subject can possess a mental capacity without exercising it—just as one can own a tool without using it. When a mental capacity is *successfully* exercised, its function is fulfilled.⁴⁵ On this line of thought, recognition is

⁴³ Cf. McGinn (1989). One issue: It seems that this epistemic strategy only explains why *we* cannot mentalistically explain recognition, not why recognition cannot be mentalistically explained *simpliciter*—as suggested by the Inexplicability Thesis.

⁴⁴ This paragraph draws on Schellenberg's (2018) discussion of mental capacities.

⁴⁵ Millar (2008) holds that the recognitional capacity can only be exercised if it is successfully exercised. I remain neutral on this point.

the exercise of a mental capacity with the function of reidentifying previously observed particulars.⁴⁶

We would like to say that some mental capacities are *more basic than* others. That is to say, at least in some cases, the possession of some capacity x depends on the possession of some other capacity y , but not vice versa.⁴⁷ It is plausible, for example, that one cannot possess the mental capacity for multiplication without possessing the capacity for addition, although to the dismay of elementary school math teachers the converse is not true. This is compatible with the idea that for some capacities x and y , the possession of x and y are mutually dependent—neither being more basic than the other. Important for our purposes, it is also compatible with the idea that some mental capacities are *bedrock* in that their possession does not depend on the possession of any other capacities. The present proposal is that the mind at its most foundational level decomposes into bedrock mental capacities.⁴⁸

If the recognitional capacity were bedrock, any attempt to provide a mentalistic explanation of recognition would be fruitless: We would not be able to mentalistically explain the capacity for recognition in any more basic mental terms. Moreover, if this suggestion were correct, the explanatory failures of Impression Matching and Recognition via Linking could be chalked up as special instances of this more general mistake.

Nonetheless, there is a *prima facie* and an *ultima facie* objection to the claim that the recognitional capacity is bedrock. The *prima facie* objection is that the recognitional capacity is too sophisticated and flexible to be a bedrock capacity. As noted in passing earlier, we are frequently able to recognize a particular despite considerable changes in its superficial appearance properties. Familiar face recognition is the paradigm example. We possess a tacit sense that a face can vary in certain ways while retaining its numerical identity. This is no small feat; after all, there is at least some sense in which a face does not look the same across changes in, say, facial expression. A picture of an individual's smiling face would be a poor pixel-by-pixel match with a picture of that same individual's grimacing face. In both children and adults, this attunement to the identity of a particular face despite such variation appears to be present only after repeated exposures to the face (Baker et al., 2017), indicating the flexibility of our capacity to recognize faces. This worry is, of course, far from decisive, given that we have no principled reason for thinking that bedrock mental capacities must be inflexible or unsophisticated.

The *ultima facie* objection is that there are at least two mental capacities whose possession is more basic than the possession of the capacity for recognition. These are the detection and discrimination capacities that function to perceptually isolate a particular from its

⁴⁶ I am *not* assuming that the only way to reidentify a previously observed particular is by recognizing it (cf. Section 2). More generally, I am not assuming that the only way to fulfil the function of a mental capacity is by exercising that capacity. There may be other ways of fulfilling that function. As an analogy, the function of a hammer is to drive in nails, but it is certainly possible to drive in nails without the use of a hammer.

⁴⁷ The relevant notion of dependence is ontological. For present purposes, it does not matter whether ontological dependence is treated as a primitive or is analyzed in terms of some other metaphysical notion, such as essence.

⁴⁸ Cf. Searle (1983).

surroundings.⁴⁹ These capacities are phylogenetically ancient. Even ants possess the capacity to discriminate harmless in-group members from dangerous out-group members, acting aggressively towards the latter but not the former (Stroeymeyt et al., 2010). I claim that one could not possess a capacity for recognition without possessing these detection and discrimination capacities.

The rationale for this last claim is that we cannot be said to possess a mental capacity if we could *never* successfully exercise it and, by extension, never fulfil its function. Applied to the present point, this means that we could not possess the capacity for recognition if we could *never* reidentify previously observed particulars. Yet, at least for creatures like ourselves, the reidentification of a particular proceeds via the successful exercise of detection and discrimination capacities that allow us to single it out from its surroundings. In order to reidentify a particular *o* at time *t*, it must be observed at *t*, and it must have been observed at some earlier time *t'*. Both the observation of *o* at *t* and the earlier observation of *o* at *t'* depend on the successful exercise of detection and discrimination capacities that allow us to single *o* out from its surroundings. If we did not possess these detection and discrimination capacities, it would not be possible for us to observe and, hence, recognize *o*. On the basis of these considerations, it appears that the recognitional capacity is not bedrock.

7. The recognitional capacity as irreducible

One might think the foregoing considerations demonstrate that an ontological answer to the meta-explanatory problem cannot be provided. This is too hasty a conclusion. For even if the possession of the capacity for recognition depends on the possession of detection and discrimination capacities, it still appears that the capacity for recognition is not *reducible* to other mental capacities. Can we make sense of this suggestion? One thing is clear: To avoid producing a question-begging answer to the meta-explanatory problem, 'reducible' must mean more than just *explainable* in this context.⁵⁰ Let us say a mental capacity is *reducible* just in case its possession fully depends on the possession of some other mental capacities; if there is no such full dependence, the capacity is *irreducible*. Thus, the capacity for recognition is reducible just in case there is some *X* such that the possession of the capacity for recognition *fully* depends on the possession of the capacity to discriminate, detect, and *X* (where *X* need not be just one capacity). If there is no such *X*, the capacity for recognition is irreducible to other mental capacities, even if it is not bedrock.

If the capacity for recognition were irreducible, any attempt to provide a mentalistic explanation of recognition would be incomplete. Such a mentalistic explanation is not merely incomplete in the sense that it is an *explanation sketch*, an outline of an explanation that one hopes can be, at least in principle, eventually filled out.⁵¹ Rather, there would always be a

⁴⁹ These capacities are discussed extensively by Schellenberg (2018). To be clear, I am interested in a (human or non-human) *subject's* capacity for detection and discrimination—that is, detection and discrimination capacities at the personal-level. Nevertheless, I do not wish to deny that these capacities can also be characterized at the subpersonal-level. The latter characterization is at play, for instance, when cognitive neuroscientists speak of edge detection in the primary visual cortex.

⁵⁰ For a discussion of explanatory and non-explanatory senses of reduction, see Sarkar (1992).

⁵¹ See Hempel (2002).

missing piece to the mentalistic explanation—the *X* that can never be included as a part of the explanans.

The final step of this ontological solution to the meta-explanatory problem involves demonstrating that the recognitional capacity is irreducible—that there is no *X* to play the part. There is at least one reason for thinking it is an irreducible mental capacity: It appears that recognition cannot be subjected to mentalistic explanation. Now, this reasoning is clearly circular, at least if the hypothesis that the recognitional capacity is irreducible is to support the claim of the Inexplicability Thesis that recognition cannot be subjected to mentalistic explanation.⁵²

Yet, there is an alternative to engaging in circular reasoning. We can treat the suggestion that the recognitional capacity is irreducible as an empirical hypothesis. Support for this hypothesis would come from evidence showing that the capacity for recognition can be possessed without possessing a range of other mental capacities (apart from detection and discrimination capacities). Although the hypothesis is clearly falsifiable, no single finding—or even collection of findings—can *verify* that it is correct; nevertheless, a large body of supporting evidence might be gradually built.

I can imagine someone objecting to this empirical approach on the grounds that it is a conceptual truth, or close to it, that *some* sort of memory—what we might call ‘recognitional’ memory—is required for recognition. But this is a red herring. For granting this truth would not establish that there is a memory capacity distinct from the recognitional capacity itself, such as an episodic or working memory capacity, that must be possessed if one is to possess a capacity for recognition. Indeed, neuropsychological evidence indicates just the opposite. The capacity for recognition is spared among elderly patients and patients with head trauma whose memory capacities are otherwise greatly impaired, as evidenced by poor performance on various explicit memory tasks (e.g., word recall tasks). Over a wide range of different delay periods, these patients’ performance on a recognition-based picture reidentification task was comparable to the performance of control subjects without any memory impairments (Spikman et al., 1995). A similar pattern of results was found among amnesia patients (Volpe et al., 1986) and in a case study involving a patient suffering from brain damage which resulted in memory loss due to meningitis (Aggleton et al., 2005).

These studies provide some indication that the possession of the capacity for recognition does not depend on the possession of any memory capacity (apart from the recognitional capacity itself). This, in turn, partially vindicates the empirical hypothesis that the recognitional capacity is irreducible.

In addition to neuropsychological evidence, there is ethological and developmental evidence for the hypothesis. A capacity for recognition is possessed by rats, as evidenced by above chance performance on object recognition tasks (Aggleton, 1985), and even certain electric fish who can recognize other individual fish by detecting the specific frequencies of their electrical organ discharges (Harvey-Girard et al., 2010). It is also well-established that human

⁵² Another problem with this reasoning is that there may be alternative explanations as to why recognition cannot be subjected to mentalistically explanation; as I mentioned above, the explanation as to why recognition cannot be mentalistically explained may have an epistemic basis.

infants possess a capacity for recognition (see Rose et al., 2004 for review). While we human adults certainly do share commonalities with rats, electric fish, and infants, these other life forms also lack many of the same mental capacities we possess, such as linguistic and certain metacognitive capacities.⁵³

Finally, there is room to combine neuropsychological, ethological, and developmental evidence to support the hypothesis. One example, which involves a hybrid neuropsychological/ethological approach, is a study showing that rats with certain impaired learning capacities—evidenced by their poor performance in learning the locations of platforms in a navigation task—had no issue recognizing familiar objects. As the experimenters put it: “In contrast to the learning of several spatial locations, object-recognition memory was normal in PDAPP animals [i.e., the group of rats in question] across their lifespan” (Chen et al., 2000, p. 977). A similar study was conducted using a group of macaque monkeys with brain damage resulting in impaired performance on a task requiring them to learn which among several items they should touch in order to receive a reward (Browning et al., 2010). Despite their poor performance on this learning task relative to a control group, the macaques’ performance was normal in an object recognition task.

These studies provide some indication that the possession of the capacity for recognition does not depend on the possession of some other learning capacity, again partially vindicating the hypothesis.

While more evidence is needed to put this empirical hypothesis on firm footing, the foregoing discussion makes it apparent that there is a non-circular support for the suggestion that the recognitional capacity is irreducible. Likewise, it is clear how to proceed. In schematic terms, we establish that there is no X such that the possession of a capacity for recognition fully depends on the possession of a capacity for discrimination, detection, and X by showing that for various candidate X s, a creature can possess a capacity for recognition (as evidenced by above-chance performance on recognition-based tasks) while failing to possess a capacity for X (as evidenced by at or below-chance performance on appropriate behavioral tasks).

8. Conclusion

I began by highlighting problems with various attempts to analyze recognition—Recognition as Re-Acquaintance, Recognition via Linking, and Impression Matching. This, in turn, led me to introduce the Inexplicability Thesis, according to which recognition cannot be subjected to mentalistic explanation. I found that the Inexplicability Thesis left us with a meta-explanatory problem: We needed an explanation for why recognition cannot be mentalistically explained. I suggested an ontological answer to this meta-explanatory problem—that the capacity for recognition cannot be mentalistically explained because it is bedrock. While this seemed to provide a tidy solution to the meta-explanatory problem, this approach turned out to be a non-starter since the recognitional capacity is not bedrock after all. Nonetheless, I argued that a

⁵³ For evidence suggesting that rats and human infants possess some metacognitive capacities, see Foote and Crystal (2007), and Goupil and Kouider (2016), respectively. To my knowledge, there is presently no evidence that fish possess any metacognitive capacities.

neighboring ontological solution is available: Recognition cannot be mentalistically explained because the recognitional capacity is irreducible.

It is important to note that the Inexplicability Thesis does not rule out various non-mentalistic explanations of recognition. As I mentioned in Section 4, we can offer neural and cognitive psychological explanations of recognition, and we can also specify its causal role with respect to other mental phenomena—what amounts to a causal explanation. One additional possibility worth mentioning, if only briefly, is that we can provide a *taxonomic* explanation of recognition. Such an explanation would enrich our understanding of recognition by identifying and demarcating its different senses. I provide a taxonomic explanation of recognition in recent work (Abid, 2021). The central distinction in this taxonomy is between the recognition of types (e.g., recognizing a dog *as such*) and the recognition of instances of types (e.g., recognizing some particular dog). In turn, familiarity and recollection are both characterized as ways of recognizing instances of types. Importantly, since this taxonomic explanation does not presuppose that there is any mentalistic explanation of recognition, it is compatible with the Inexplicability Thesis.

The distinction between these two different senses of recognition brings an outstanding issue into focus. While we have concentrated on the recognition of *instances* of types, one wonders whether parallel considerations apply to the recognition of types. For Evans, the answer would seem to be yes: “For of course what I say about the process of recognition applies no less to the recognizing of something as a dog (the recognition of a type) than to the recognition of particular instances of that type” (1982, p. 289). Unfortunately, a discussion of this last point will have to wait until another time.

In the next and final chapter, I turn to a widespread disruption of our recognitional capacity that has serious practical consequences: the cross-race effect. I argue that the cross-race effect has been mischaracterized by contemporary psychologists. By more accurately characterizing the cross-race effect, we put ourselves in a better position to help mitigate some of its harmful effects.

Chapter 4: Towards a two-factor approach to the cross-race effect

1. Introduction

The cross-race effect (*CRE*) is standardly characterized as the finding that individuals are generally better at recognizing previously observed faces of members of their own race (*own-race faces*) than faces of members of other races (*other-race faces*).⁵⁴ The CRE has been replicated extensively (Meissner and Brigham, 2001; Young et al., 2012). Likewise, it has been observed in individuals belonging to different racial groups (Ng and Lindsay, 1994) and in children less than a year old (Kelly et al., 2007).

While the existence of the CRE is largely undisputed, there is little agreement about the mechanisms underlying it. Two theories have dominated. *Perceptual expertise* theories claim the CRE is a perceptual effect that derives from a lack of contact with other-race faces. *Social-cognitive* theories claim the CRE is explained not by a lack of perceptual expertise, but by a difference in ingroup/outgroup classification that causes an other-race face to be processed differently from an own-race face. While these two theories are often framed as competitors, hybrid theories have combined elements of both (see Young et al., 2012).⁵⁵

My central thesis is that the CRE has been mischaracterized in the contemporary psychological literature. It is widely assumed that there is just one factor associated with the CRE: a difficulty in recognizing other-race faces relative to own-race faces (*Factor 1*). This much is assumed by perceptual expertise and social-cognitive theories (discussed in more detail in Sections 2-3), along with their corresponding hybrid theories. I argue that this assumption is flawed (Section 4). In addition to Factor 1, the CRE involves a limited metacognitive awareness of one's difficulty in recognizing other-race faces (*Factor 2*). In contrast to existing *one-factor* approaches that focus exclusively on Factor 1, I advocate a *two-factor* approach to the CRE that seeks to explain both of its associated factors.

In the second half of the chapter, I focus on developing a specific version of the two-factor approach. In general, a two-factor approach can draw on existing theories of the CRE to account for Factor 1 but must provide a supplementary explanation of Factor 2. I hypothesize that Factor 2 is a product of our tacit yet mistaken assumption that we can recognize individuals of all races equally well (Section 5). One prediction of this hypothesis, which is borne out, is that subjects *judge* themselves to be equally likely to recognize individuals regardless of their race. This prediction is not offered by a competing hypothesis which appeals to our general tendency towards overconfidence when task performance is poor. I show how my preferred two-factor approach can be extended to apply to neighboring phenomena, such as our difficulty in recognizing the faces of members of age groups different than our own (Section 6).

There are several motivations underlying the development of a two-factor approach. The first concerns the practical consequences of the CRE. One particularly grave consequence pertains to eyewitness testimony. It has been estimated that over 1/3 of wrongful convictions in

⁵⁴ The cross-race effect is also called the “own-race bias” or the “other-race effect.”

⁵⁵ In emphasizing that the CRE involves both perceptual deficits (as emphasized by perceptual expertise theories) as well as cognitive deficits (as emphasized by social-cognitive theories), hybrid theories align with the view that recognition constitutively involves *both* perception and cognition (Abid, 2022).

the United States involve the misidentification of an individual belonging to a race other than the witness's (Innocence Project, 2008). The CRE is a likely cause of many of these eyewitness misidentifications and, by extension, wrongful convictions. The CRE can be dehumanizing in other ways as well. For example, by frequently misidentifying individuals belonging to other races, we (as teachers, supervisors, etc.) fail to adequately acknowledge them and thereby create an unwelcoming atmosphere. The fact that this consequence is unintended does not curtail its effects. In due course, I argue that a two-factor approach to the CRE helps us better understand the harms it causes.

A second motive concerns the connection between the CRE and racism. This connection can be investigated through a series of related questions. Is the CRE a cause of racism? Is the CRE caused by racism? Is CRE itself racist? Answers to these questions turn on our understanding of the nature of both racism and the CRE. For example, we would expect different answers to these questions depending on whether individuals or social institutions are assumed to be conceptually foundational in understanding racism. Similarly, our answers will vary depending on how exactly the CRE is characterized. Part of the motivation of the present chapter is to sharpen our understanding of the CRE so that we are in a better position to provide such answers.

A final motivation for developing a two-factor approach is to pinpoint why exactly the CRE is epistemically problematic, a task more difficult than one might initially expect. Assuming that our face recognition capacities must be 'tuned' to certain groups of faces at the expense of others, it would seem to make sense from an epistemic perspective to tune our recognitional capacities to own-race faces as opposed to other-race faces, given that we tend to interact with members of our own race more often than members of other races. Assuming this tuning strategy maximizes the overall reliability of our recognitional capacities, it becomes hard to say what is epistemically defective about our enhanced ability to recognize own-race faces. I argue that we can identify what is epistemically problematic about the CRE by instead focusing on the metacognitive limitation associated with it, which systematically distort our understanding of our own recognitional capacities (an issue I return to in Section 7).

2. Perceptual expertise theories

In this section and the next, I provide a brief overview of existing psychological theories of the CRE. I focus primarily on perceptual expertise and social-cognitive theories, beginning with the former.

Perceptual expertise theories start with the observation that people are often, though not always, exposed to individuals of their own race much more frequently than individuals of other races. This frequent exposure is said to result in a perceptual expertise for own-race faces that facilitates recognition of these faces relative to other-race faces. In turn, the CRE is explained in terms of this selective perceptual expertise.

Perceptual expertise theories are supported by evidence that increased contact with other races can reduce and even reverse the CRE. For instance, Chiroro and Valentine (1995) find that the CRE for White faces is attenuated in Black individuals with a high degree of contact with Whites. Similarly, Sangrigoli et al. (2005) show that Korean children adopted between the ages of 3-9 *years* and raised in France recognize White faces more accurately than Asian faces. Since

children first exhibit the CRE by 6-9 *months* of age (Kelly et al., 2007), long before any of the children in this study were adopted, we can be confident that this was a genuine reversal of the standard CRE.

What kind of processing advantage would perceptual expertise confer? The answer varies depending on the version of the perceptual expertise theory in question. Michel et al. (2006) stress that perceptual expertise allows for an own-race face to be processed in a configural fashion; in contrast, other-race faces are processed in a feature-based, or piecemeal, fashion. Although exact definitions of “configural” and “feature-based” processing differ among experimenters, the rough idea is that relational information about facial landmarks (e.g., nose-to-mouth distance) is encoded in configural processing, whereas intrinsic information about facial landmarks (e.g., nose shape) is encoded in feature-based processing. The tacit assumption is that configural processing is more effective for familiar face recognition than feature-based processing.⁵⁶ Papesch and Goldinger (2010) hypothesize that the CRE arises due to a difference in how own- and other-race faces are clustered in “face space,” a n-dimensional coordinate system where each dimension corresponds to some physiognomic aspect and each individual’s face is represented as a point. The suggestion is that own-race faces are represented sparsely in face space as a result of increased perceptual expertise, facilitating differentiation between own-race faces; in contrast, other-race faces are represented in dense clusters in face space that makes differentiation more difficult.

A challenge faced by perceptual expertise theories is that the evidence regarding the correlation between increased exposure to other races and the CRE is mixed. While some studies, such as those cited above, find a negative correlation between the CRE and other-race exposure, others do not (e.g., Ng and Lindsay, 1994). Similarly, other-race exposure seems to account for only a small amount of the variability in the CRE, roughly 2% (Meissner and Brigham, 2001).

3. Social-cognitive theories

According to social-cognitive theories, when an individual’s face is perceived, it is classified as belonging to either an ingroup or outgroup member. If classified as belonging to an ingroup member, various mental resources, such as attentional, memory, and motivational resources, are used to *individuate* that that face as effectively as possible, thereby facilitating later recognition of that face. If classified as belonging to an outgroup member, these resources are instead used to *categorize* the face in terms of (e.g.) race, gender, or age. Since other-race faces tend to be classified as belonging to an outgroup member, resources are dedicated to categorizing these faces, not to individuating them.⁵⁷ This pattern is reversed for own-race faces, which tend to be classified as belonging to an ingroup member. Accordingly, this difference in ingroup/outgroup classification and mental resource allocation is what gives rise to the CRE.

There are different versions of social-cognitive theories. On Levin’s (2000) version, when a face is classified as belonging to an ingroup member, a subject encodes facial features

⁵⁶ For a critique of this assumption, see Burton et al. (2015).

⁵⁷ For the sake of simplicity, I treat individuation and categorization as if they were all-or-nothing, binary options. In reality, there is likely a graded spectrum of options, with ‘pure’ forms of individuation and categorization sitting on opposite extremes. None of my argumentative points turn on this simplification.

that are most helpful for individuating that face. When a face is classified as belonging to an outgroup member, encoding is restricted to those facial features most useful for categorization. On Rodin's (1987) version, outgroup faces are altogether "cognitively disregarded." On this view, ample mental resources are allocated to the processing of a face only if it belongs to an ingroup member. For faces classified as belonging to an outgroup member, resources are only allocated to allow for superficial face processing.

Social-cognitive theories offer a different set of predictions about the CRE than perceptual expertise theories. For instance, in the study of Sangrigoli et al. (2005) discussed above, social-cognitive theories predict that adopted Korean children would exhibit the usual CRE if the outgroup status of the surrounding White Europeans was made salient. Likewise, social-cognitive theories predict that CRE-like impairments can be generated on the fly, without any differences in perceptual expertise, as long as there is a difference in ingroup/outgroup classification. This prediction receives some vindication. Artificial ingroup/outgroup distinctions created in a laboratory suffice to generate CRE-like impairments, even when prior exposure to these ingroup/outgroup members is the same (Bernstein et al., 2007).

The aforementioned studies provide some support for social-cognitive theories. However, other studies paint a different picture. For example, Rhodes et al. (2010) find that manipulations affecting whether a racially ambiguous face is classified as either White or Asian had *no* influence on White participants' ability to later recognize the face. This finding is at odds with the prediction offered by social-cognitive theories. Thus, just as is the case with perceptual expertise theories, the evidence for social-cognitive theories is on the whole mixed. In the next section, I argue that equivocal evidence is not the only problem faced by these theories.

4. The problem with one-factor approaches

At a certain level of generality, perceptual expertise and social-cognitive theories—along with corresponding hybrid theories, which I have only mentioned briefly—both adopt a similar explanation of the CRE. Both are examples of what I call *one-factor* approaches. On a one-factor approach, the CRE arises because own-race faces tend to be processed in a way that is more advantageous for recognition relative to other-race faces. According to perceptual expertise theories, this recognition advantage arises because of a selective perceptual expertise for own-race faces. According to social-cognitive theories, the advantage arises because of the ingroup classification of and consequent mental resource allocation for own-race faces.

One-factor approaches clearly explain a core aspect of the CRE. It is difficult to see how we could acknowledge the existence of the CRE without admitting this much. However, I contend that one-factor approaches do not provide a *complete* explanation of the CRE. Specifically, they do not explain why subjects are overconfident in their capacity to recognize other-race faces—an effect indicative of a limited metacognitive awareness of their difficulty in recognizing such faces. If this is right, then existing one-factor approaches are in need of supplementation.

One of the clearest studies providing evidence for this overconfidence effect is that of Dodson and Dobolyi (2016). In this study, experimenters ask participants to recognize own- and other-race faces presented during an earlier encoding phase in a mock lineup, which consists of images of six different faces. Some lineups include only own-race faces, while others include

only other-race faces. In some lineups, all six faces have never been seen before by participants; in other lineups, one of the six faces has been previously seen. For each lineup, participants are required to indicate whether they recognize any of the six faces and provide a confidence rating about the accuracy of their response. In addition to replicating the standard CRE, Dodson and Dobolyi find that confidence ratings are moderately well calibrated with task accuracy for own-race faces.⁵⁸ For other-race faces, however, calibration is significantly worse. Importantly, this is due to participants' overconfidence in recognizing other-race faces. One-factor approaches do not account for this finding—and others like it.⁵⁹ As a result, they fail to address the metacognitive aspect of the CRE.

This objection to one-factor approaches presupposes that a metacognitive deficit partly constitutes the CRE. But why should we accept this? Couldn't we instead acknowledge the existence of this metacognitive deficit while nonetheless denying that it is *definitive* of the CRE? Such a response would align with the way contemporary psychologists characterize the CRE and would suggest that one-factor approaches are on the right track after all.

In response, it is worth clarifying that my aim in characterizing the CRE is revisionary, not descriptive. That is, my aim is to clarify how the CRE *should* be characterized, not how it is *in fact* characterized by theorists. Thus, it is no strike against my approach that it deviates from the orthodox construal of the CRE. Moreover, my guiding assumption is that the CRE ought to be defined and redefined in such a way that best suits our needs as theorists with the goal of understanding it and bringing attention to its harms.⁶⁰ By endorsing the extant definition, which excludes the accompanying metacognitive deficit, we make it easier to neglect the harms associated with the CRE. In contrast, we highlight these harms by broadening the existing definition of the CRE so as to include its accompanying metacognitive deficit.

Concrete examples help to bring this last point into focus. In light of findings concerning the CRE, the Innocence Project (2020), which works to combat wrongful convictions, recommends that eyewitnesses provide confidence ratings about their identifications as a means of combating wrongful misidentification. In making this recommendation, the Innocence Project is operating with a narrow definition of the CRE as involving just one factor. Given this definition, the proposed recommendation seems unproblematic and potentially useful. Nevertheless, it is in fact antithetical to the aim of combating witness misidentification since an overconfident but mistaken eyewitness is more convincing than a mistaken eyewitness who provides no indication of their confidence. By contrast, the broader definition that I favor makes it apparent that a subject's reported confidence is not a reliable means of assessing how accurate they are in recognizing the face of an individual of another race. Eyewitnesses testifying against

⁵⁸ Note that calibration is still imperfect in the case of own-race recognition, reflecting a smaller degree of overconfidence than in the case of other-race recognition. *Perfect* calibration would be surprising given the general pattern that we have only limited insight into our face recognition capacities (see Palmero et al., 2017).

⁵⁹ Hourihan et al. (2012) provide evidence of a metacognitive deficit by showing that participants are less likely to accurately predict their performance on a face recognition task involving an other-race face than an own-race face. Unfortunately, the experimenters do not report whether this effect is driven by participants' overconfidence. I outline additional evidence for this overconfidence effect in 5.3.

⁶⁰ This assumption is closely aligned with what Haslanger (2000) calls an *analytic approach*, according to which our answers to questions of the form "What is *X*?" are informed by a consideration of what we want the concept *X* to do for us. On an analytic approach, concepts are defined partly in relation to our aims.

individuals of other races are liable to make *confident* misidentifications. These issues remain opaque under the narrow definition of the CRE employed by most psychologists.

Similarly, the broader definition of the CRE I favor makes it clear that you and I are almost certainly susceptible to the CRE *even if we do not think we are*. Indeed, on my favored definition, it is part of the very nature of the CRE to think that we are not susceptible to it. By contrast, the narrow definition of the CRE endorsed by most psychologists provides no reason to second-guess our own metacognitive assessment. This, in turn, makes it easier to neglect the everyday harms we cause as a result of our susceptibility to the CRE. For example, it makes it easier to overlook that we regularly alienate members of other races by failing to recognize them in social situations.

These examples illustrate how my favored definition of the CRE more effectively highlights its harms. However, even if one is not moved by the consideration that a definition of the CRE should bring attention to its deleterious consequences, there is still reason to favor my broader definition on the grounds that it sharpens our understanding of the CRE. To appreciate this point, let us consider a proposal by Malinowska (2016) that the CRE is an instance of a more general phenomenon known as the unfamiliar homogeneity effect (UHE). Unlike the CRE, the UHE does not only involve faces. It also includes "...similar difficulties relating to other social groups, languages, sounds and objects, whose recognition we are not specialised in" (ibid., pp. 3866-7). For example, people, such as myself, who rarely listen to classical music often say it all sounds the same. Thus, I exhibit a UHE for classical music. Similarly, those of us who do not spend extensive amounts of time around pigeons have trouble recognizing individual pigeons and therefore exhibit a UHE for pigeons.

Malinowska does not discuss the issue of metacognitive awareness in light of the UHE. Still, it is reasonable to ask whether we are overconfident in our capacity to recognize those various types of things we do not specialize in. Is it the case, for example, that I am not only bad at distinguishing Beethoven from Bach but am also overconfident in my ability to do so? Do I take myself to be more skilled in identifying individual pigeons than I actually am? Although these are ultimately empirical questions, there is plausibly an asymmetry between these cases and the case of human face recognition. By my own lights, I am not overconfident in my capacity to recognize the works of classical composers; instead, I readily acknowledge my limitations in this domain.⁶¹

If this line of thought is right, we should resist the temptation to regard the CRE as a mere instance of the UHE. This suggestion is tempting only if one adopts a one-factor approach, as Malinowska does.⁶² However, by adopting a broader definition of the CRE that acknowledges

⁶¹ This last assessment might be a product of a *meta-metacognitive* impairment: a poor understanding of my own metacognition. However, such a meta-metacognitive impairment would not account for my hunch that *you* are not susceptible to overconfidence in *your* ability to recognize individual pigeons.

⁶² Malinowska claims that the UHE emerges because of a lack of perceptual expertise (or "perceptual narrowing") with certain social groups, objects, faces, etc. Putting two and two together, it is apparent that Malinowska is ultimately advocating a perceptual expertise theory, albeit one that has a broader scope than standard perceptual expertise theories which focus on just the CRE. As such, Malinowska is adopting a one-factor explanation of the CRE.

its accompanying metacognitive deficit, we acknowledge the possibility that it is importantly different from other examples of the UHE that do not involve the recognition of human faces.

Let me highlight a final concern with adopting a two-factor approach to the CRE. Consider the separate phenomenon of implicit bias. Like the CRE, implicit bias involves both a biased response that is manifested behaviorally and a lack of metacognitive awareness of that bias. Yet, theorists in the implicit bias literature usually take themselves to be explaining just one thing (viz., implicit bias), not two (viz., bias and a lack of metacognitive awareness of it). Hence, it would be odd to suggest that a two-factor explanation of implicit bias is needed. So, why is there a need for a two-factor approach in the case of the CRE? Where does the disanalogy lie?⁶³

To appreciate the disanalogy, note that implicit bias is standardly characterized as involving limited metacognitive insight. For example, the American Psychological Association defines implicit bias as “a negative attitude, *of which one is not consciously aware*, against a specific social group” (2023, italics added). A lack of awareness of one’s own biased attitudes is commonly taken to be an essential and defining feature of implicit bias and is *already* incorporated into existing accounts of the phenomenon. Thus, a two-factor approach to implicit bias would be redundant. By contrast, a lack of metacognitive awareness is not a part of the standard characterization of the cross-race effect. For example, the perceptual expertise, social-cognitive, and hybrid theories introduced in Sections 1-3 do not characterize the cross-race effect as involving any metacognitive deficit. If the cross-race effect *were* standardly characterized as involving a limited metacognitive awareness, as in the case of implicit bias, then a two-factor approach would be unnecessary and existing accounts would not require supplementation. However, that is not the situation we are in.

To briefly recap, I have been arguing that we ought to revise our understanding of the CRE and accept that two factors constitute it: a difficulty in recognizing other-race faces relative to own-race faces coupled with a limited awareness of this difficulty—an evaluative limitation in acknowledging one’s own difficulty in recognizing other-race faces. Again, my reason for this endorsing this revisionary characterization is that it helps us better understand the CRE (by distinguishing it from the UHE) and more effectively highlights its harms (as illustrated by the work of the Innocence Project and our own susceptibility to the effect). What is needed is a *two-factor* approach which seeks to explain both factors associated with the CRE, in contrast to a one-factor approach which accounts for only the former.⁶⁴ Finally, parallel considerations do not apply in the case of implicit bias since a metacognitive deficit is already part of the standard characterization of implicit bias and incorporated into existing accounts of the phenomenon.

5. Towards a two-factor approach

A two-factor approach can draw on a perceptual expertise, social-cognitive, or hybrid theory to explain the first factor associated with the CRE—our difficulty in recognizing other-race faces relative to own-race faces. But what would explain the second factor associated with the CRE—our limited awareness of this difficulty? Simply citing our limited insight into face recognition is not enough: Such an answer does explain why we possess a greater metacognitive awareness of

⁶³ I thank an anonymous reviewer for raising this concern.

⁶⁴ To my knowledge, Rhodes et al. (2013) are the only theorists who explicitly suggest that the CRE involves a metacognitive impairment. However, they do not emphasize the need for a two-factor approach.

our capacity to recognize own-race faces in comparison to other-race faces. In this section, I explore two candidate explanations for Factor 2 (5.1 and 5.2) and argue that we have reason to favor the latter on the grounds that it offers a more fine-grained prediction regarding subjects' performance on face recognition tasks (5.3).

5.1 A generic two-factor approach

I begin with what I call a *generic two-factor approach*. On this approach, Factor 2 is an instance of a generic metacognitive deficit. By a *generic metacognitive deficit*, I have in mind one that applies to individuals' assessments of their own performance regardless of which exact task they are performing. We can think of a generic metacognitive deficit as one that is task insensitive and domain general.

The most promising way to develop the generic two-factor approach is by appeal to the *Dunning-Kruger effect*, our tendency to overestimate our performance when we are least proficient and underestimate it when we are most proficient (see Kruger and Dunning, 1999).⁶⁵ An example of this effect: Students who score in the top quartile on an exam tend to underestimate their performance, whereas students in the bottom quartile tend to overestimate it (Ehrlinger et al., 2008).⁶⁶ Importantly, there is nothing special about students' exam performance. The effect is also exhibited by bridge players, physicists, and pharmacists (see Dunning, 2011 for review). Indeed, the effect even extends to the domain of face perception (Zhou and Jenkins, 2020). This is unsurprising given that the Dunning-Kruger effect is a generic metacognitive deficit (in the aforementioned sense).

The Dunning-Kruger effect offers a possible explanation for the metacognitive deficit associated with the CRE. In light of our susceptibility to the Dunning Kruger effect, we have a tendency to overestimate our performance on *any* task where our performance is poor. The recognition of faces of members of other races is one such task. (Again, we can appeal to existing theories of the CRE to explain why this is so.) Thus, it is expected that we would overestimate our performance on tasks requiring us to recognize other-race faces.

5.2 A non-generic two-factor approach

A *non-generic two-factor approach* offers a contrasting explanation of Factor 2. This approach holds that the metacognitive deficit associated with the CRE is *non-generic* in the sense that it applies to individuals' assessments of their own performance on only a circumscribed range of

⁶⁵ The Dunning-Kruger effect is a psychological generalization. Like psychological generalizations in general, it admits of occasional counterexamples. Thus, an extremely modest student who scores in the bottom quartile on an exam might nonetheless *underestimate* their performance. It is also worth pointing out that who lack even the most basic proficiency on a given task do not exhibit the effect. As Kruger and Dunning themselves note: "...most people have no trouble identifying their inability to translate Slovenian proverbs, reconstruct an 8-cylinder engine, or diagnose acute disseminated encephalomyelitis. In these domains, without even an intuition of how to respond, people do not overestimate" (1999, p. 1132). This might explain why, for example, those who lack even a rudimentary capacity to recognize individual pigeons are not overconfident in their capacity to do so (cf. Section 4).

⁶⁶ This holds true of both absolute and relative performance estimates. For instance, a student in the bottom quartile with overestimate their raw score and will also overestimate their performance relative to other students (see Ehrlinger et al., 2008). Interestingly, overestimation on exam performance is most pronounced among male students (Bengtsson, Persson, and Willenhag, 2005).

tasks. In contrast to a generic metacognitive deficit, which is task insensitive and domain general, a non-generic deficit is task sensitive and domain specific.

Although a non-generic two-factor approach might be developed in several different ways, I focus on one specific version of the approach in which the metacognitive deficit associated with the cross-race effect is explained by our tacit endorsement of an assumption that we recognize faces of all races equally well. I call this the *homogeneity* assumption. Importantly, the homogeneity assumption is flawed: We do not recognize faces of all races equally well. Thus, when paired with a difficulty in recognizing other-race faces relative to own-race faces, a tacit endorsement of the homogeneity assumption leads us to overestimate our ability to recognize other-race faces.

I understand tacit endorsement as a type of implicit attitude. Like implicit attitudes in general, a tacit endorsement can guide an individual's behavior without cohering with their explicit beliefs. For instance, Dovidio and Gaertner (2000) find that while there was a decline in White subjects' self-reported racial prejudice towards Blacks between 1989 and 1999, there was no corresponding decline when more subtle measures of racial prejudice were used. A standard interpretation of this finding is that while White subjects' avowed beliefs about Blacks were revised during the 1989 to 1999 period, their corresponding implicit attitudes remained largely unchanged. In this example, there is a mismatch between the content of subjects' explicit beliefs and their implicit attitudes. However, implicit attitudes can also fail to cohere with explicit beliefs not because of inconsistency but because a subject lacks any corresponding explicit beliefs. The latter is relevant for our purposes. In claiming that subjects tacitly endorse the homogeneity assumption, I am not claiming that subjects explicitly believe the homogeneity assumption, nor am I claiming that they explicitly reject it. I would wager that most subjects have no considered belief whatsoever concerning their ability to recognize faces belonging to individuals of other races. This vacuum in explicit belief opens the possibility that different subjects will respond in different ways when presented with evidence that they themselves are susceptible to the CRE. Some may vehemently deny that they have any difficulty recognizing members of other races—perhaps out of shame, guilt, or fear that such an admission may implicate them as racists. By contrast, those with greater humility might openly concede their susceptibility to the effect once it is pointed out.

On the present proposal, the homogeneity assumption is the *content* of one's tacit endorsement. Moreover, it is a *propositionally structured* content, roughly, a truth-evaluable content with a structure reminiscent of a natural language sentence. Since I regard tacit endorsement as an implicit attitude, it follows that the tacit endorsement of the homogeneity assumption is an implicit attitude with a propositionally structured content. This commitment is compatible with theories that characterize implicit attitudes as fragmented, unconscious beliefs (Mandelbaum, 2016). It is also compatible with theories which regard implicit attitudes as *sui generis* mental states that are similar to beliefs in that they possess propositionally structured contents (Levy, 2015). However, it is *incompatible* with views according to which implicit attitudes are *sui generis* mental states with “associative” contents that are not propositionally structured (Gendler, 2008). While some implicit attitudes may have associative contents, if the present proposal is correct, not all do.

One substantive question is why anyone would tacitly endorse the homogeneity assumption, given that it is incorrect. One possibility is that it is endorsed by default—perhaps due to its simplicity and parsimony in comparison to the contrasting assumption that our face recognition capacities are sensitive to race. Moreover, this default assumption would never be disconfirmed by most subjects. There are several reasons for this. First, we rarely, if ever, receive extensive feedback concerning our ability to recognize other-race faces outside of laboratory settings. Similarly, while the CRE is well-replicated, a large portion of society is nonetheless ignorant of it. Without any such feedback or some general knowledge of the CRE, a subject’s default tacit endorsement of the homogeneity assumption would remain unchecked. Furthermore, we lack introspective access to the mechanisms which give rise to the first factor associated with the CRE. For example, on the approach of Michel et al. (2006), subjects’ difficulty in recognizing other-race faces relative to own-race faces is caused by the use of feature-based as opposed to configural processing (Section 2). The “decision” to use feature-based processing—if we wish to call it that—is made not by the subject, but by the information processing systems *within* the subject. More generally, we lack introspective access to the inner workings of these systems. Since the processing difference between own- and other-race faces remains introspectively opaque to us, it does not occur to us that we process own- and other-race faces differently. Consequently, the homogeneity assumption goes unchallenged.

5.3 Deciding between generic and non-generic two-factor approaches

I have sketched two ways of developing the two-factor approach. The first is a generic approach that explains Factor 2 by appeal to the Dunning-Kruger effect; the second is a non-generic approach that explains Factor 2 by appealing to our tacit endorsement of the homogeneity assumption. How do we decide between these two approaches?

We can begin to answer this question by considering the predictions of each two-factor approach. Both approaches share a prediction. Like the generic approach, the non-generic approach predicts that subjects will overestimate their ability to recognize other-race faces. However, unlike its generic counterpart, the non-generic approach offers a more fine-grained prediction concerning subjects’ responses on face recognition tasks. It predicts that this overestimation will occur in a highly specific way: If subjects are tacitly endorsing the homogeneity assumption, they will judge themselves *equally* likely to recognize other-race faces as own-race faces.

The upshot is that the two approaches can be empirically disentangled in light of their differing predictions. Furthermore, these differing predictions can help us decide between the two approaches. If the fine-grained prediction of the non-generic approach is correct, we would have reason to favor it over the generic approach. On the other hand, if this prediction is incorrect, that would count as a strike against the non-generic approach, and we would have reason to instead favor the generic approach.⁶⁷

⁶⁷ It is worth making the caveat that this fine-grained prediction is offered by the specific version of the non-generic approach I have developed, that is, the non-generic approach which involves a tacit endorsement of the homogeneity assumption. I am not claiming that every version of the non-generic approach would make this fine-grained prediction. I am only concerned with the predictions of my favored version of the non-generic approach. Similarly, I do not wish to deny that some version of the generic approach could make such a prediction. I am only claiming that

Is this fine-grained prediction borne out? Although evidence is still forthcoming, a study by Rhodes et al. (2013) suggests an affirmative answer. In their first experiment, Rhodes and colleagues ask White participants to study Black and White faces presented sequentially for two seconds each. Participants are told that they will be required to recognize these faces later. Following each face presentation, participants provide a “judgment of learning,” a scaled rating indicating their predicted likelihood of later recognizing that face. If participants were aware of their difficulty in recognizing other-race faces relative to own-race faces, their judgments of learning would be lower for other-race faces than for own-race faces. However, Rhodes and colleagues find that participants’ judgments of learning do not reliably differ for Black and White faces, even though they exhibit the standard difficulty in recognizing Black faces compared to White faces during the later recognition task. This is exactly what would be expected if participants were tacitly endorsing the homogeneity assumption.

One potential pitfall of this first experiment is that the aims of the experimenters may be transparent to participants. It is therefore possible that participants are aware of their difficulty in recognizing other-race faces relative to own-race faces but try to conceal this difficulty by reporting similar judgments of learning for own- and other-race faces in an effort not to embarrass themselves in front of experimenters.

This concern is addressed by a follow-up experiment which uses a more subtle measure of participants’ metacognitive awareness: study time. In their second experiment, Rhodes and colleagues make one small adjustment to the first experiment. Instead of using a fixed presentation time, they now allow participants decide how long they wish to study each face during the initial presentation stage. Unbeknownst to participants, the experimenters record the amount of time they devote to studying the own- and other-race faces they are required to recognize at a later point in time. If participants were attuned to their own deficit in recognizing other-race faces, they would have devoted a longer amount of time studying other-race faces. What Rhodes and colleagues actually find is that participants spend roughly equal amounts of time studying own- and other-race faces.⁶⁸ Again, this is just what would be expected if participants were tacitly endorsing the homogeneity assumption.

6. The challenge of neighboring phenomena

In this section, I consider a challenge to the non-generic two-factor approach I developed in Section 5. The trouble is that we find effects similar to the CRE when it comes to recognizing faces of members of other groups. For instance, we tend to have difficulty recognizing the faces of members of other *age* groups relative to those of members of our own age group (e.g., a teenager generally has more difficulty recognizing the face of an adult than another teenager). This phenomenon is known as the *cross-age effect* (Rhodes and Anastasi, 2012). Just as in the case of the CRE, we have a limited metacognitive awareness of this difficulty. For example, young adults who have difficulty recognizing older faces are overconfident in their capacity to do so (Bryce and Dodson, 2013).

the version of the generic approach under consideration—what I take to be the most promising version involving the Dunning-Kruger effect—does not make this prediction.

⁶⁸ A similar pattern of results is obtained by Tullis et al. (2014).

Given the similarity of the cross-age effect to the cross-race effect, you might think that they should both be explained in the same way. That is, you might think an explanation of the one should subsume the other. However, the non-generic approach that I've sketched in Section 5 does not explain the metacognitive deficit associated with the cross-age effect. This is because the non-generic approach explains the metacognitive deficit in terms of the homogeneity assumption, which pertains exclusively to race.

Now, there is an easy fix, which allows the non-generic approach to accommodate the metacognitive deficit associated with the cross-age effect. The idea would be to revise the homogeneity assumption by making it more encompassing. The assumption would now be that we recognize faces of all *groups* equally well, not just faces of all races. A tacit endorsement of *this* assumption leads to overconfidence in recognizing the faces of members of other groups (e.g., other racial and age groups) when we have more difficulty recognizing them.

One wonders whether this revision salvages the non-generic two-factor approach in letter, but not in spirit. By making the homogeneity assumption more encompassing, aren't we admitting that there is something right about a generic two-factor approach after all? If this is correct, then the revision is in a way self-undermining to the non-generic approach. I think this is the deeper worry that is highlighted by neighboring phenomena such as the cross-age effect.

In response, it is important to appreciate that the distinction between generic and non-generic two-factor approaches is fuzzier than we initially supposed. Reflection on the cross-age effect forces us to concede that the metacognitive deficit associated with the CRE generalizes beyond just race. At the same time, we still have reason to doubt that the metacognitive deficit is a generic one. After all, this suggestion does not generate the fine-grained prediction offered by the non-generic approach sketched in Section 5. We are left with a two-factor approach that cannot be straightforwardly characterized as either generic or non-generic; rather, it sits somewhere along a generic/non-generic continuum.

7. Conclusion

In sum, existing theories of the CRE fail to explain subjects' limited awareness of their own difficulty in recognizing faces of other-races. To address this worry, I recommended that we move towards a two-factor approach to the CRE. I sketched one specific version of a non-generic two-factor approach that explains subjects' limited metacognitive awareness in terms of a faulty assumption that they recognize faces of all races equally well, arguing that it offers more precise predictions than a generic two-factor approach which explains subjects' limited awareness in terms of in terms of a task insensitive and domain general metacognitive deficit, namely, the Dunning-Kruger effect. Finally, I showed how my favored two-factor approach can be revised to accommodate neighboring phenomena, such as the cross-age effect. The cost of this revision is that my two-factor approach can no longer be easily characterized as a *non-generic* two-factor approach.

To conclude, I would like to highlight the implications of adopting a two-factor approach on our understanding of the epistemology of the CRE. It is difficult to pinpoint what is epistemically bad about the CRE if one endorses a one-factor approach in which the CRE consists entirely in a difficulty in recognizing members of other races. Munton (2019) asks whether the development of an advantage in recognizing own-race faces over other-race faces

could in fact be epistemically beneficial. The scenario Munton has in mind is one in which our recognitional capacities must be tuned to certain groups at the expense of others. In this scenario, face recognition is a zero-sum game: Aptitude in recognizing faces belonging to one group is always balanced by ineptitude in recognizing faces belonging to another group. While we might question whether some (or all) of our recognitional capacities work this way, let us bracket this question for now and suppose that such a tradeoff is necessary. If this is right, it would make sense from an epistemic perspective to tune the recognitional capacity to faces of those groups of individuals we most commonly interact with. Given the persisting racial segregation in many countries, including the United States, one consequence would be that our face recognition capacities would often be tuned to own-race faces at the expense of other-race faces.

What would be epistemically wrong with this tuning strategy? Given the nature of one's surroundings, tuning the recognitional capacity in this way would seem to optimize its overall reliability, even if it yields occasional false beliefs regarding the identity of members of other races.⁶⁹ While the strategy clearly has harmful ethical and practical consequences, it is difficult to point out why it is flawed on epistemic grounds. This makes it difficult to say what, if anything, is epistemically problematic about our difficulty in recognizing other-race faces.

Munton (forthcoming) suggests that the problem with optimizing our recognitional capacities in this way is that we limit our ability to respond to faces in a demographically neutral fashion. In other words, by adopting this tuning strategy, we are no longer in a position to disregard race in the process of recognizing a familiar face. Munton suggests that there is epistemic value in “a capacity to learn about others in ways that disregard their demographic status” (ibid.). This view faces some challenges. *Pace* Munton, it is difficult to come up with examples illustrating the epistemic value of demographic neutrality that cannot be explained away in terms of its practical or moral value. Of course, there are many examples in which subjects make faulty inferences on the basis of demographic information, such as when it is inferred that differences in IQ scores across racial lines are explained by the genetic inferiority of certain races (see Block, 1996 for critical discussion), but in such cases it is the faulty inference that is epistemically blameworthy, not the consideration of demographic information *per se*. More generally, it is unclear why disregarding *any* information—demographic or otherwise—should be epistemically valuable. If anything, one might think that demographic neutrality leaves us in an epistemically disadvantageous position, even if it has clear practical benefits (e.g., making us less likely to behave in biased ways).

By adopting a two-factor approach, we can nonetheless appreciate what is epistemically problematic about the CRE in a way that bypasses these worries. For even if we accept that there is nothing epistemically wrong with our difficulty in recognizing members of other races (Factor 1), it is epistemically problematic that we fail to acknowledge our difficulty in recognizing other-race faces relative to own-race faces (Factor 2) and proceed as if the limitation does not exist.⁷⁰ This, in turn, reflects a lack of self-knowledge concerning the limits of our own capacities. In

⁶⁹ This point can be appreciated even if one rejects a reliabilist approach to epistemic justification. For a classic discussion of reliabilism, see Goldman (1979).

⁷⁰ Note that Factor 2 might offer certain adaptive or even practical benefits, even if it is epistemically problematic. For example, a creature with a heightened metacognitive awareness of all its limitations might get distracted by some of them. This is something that evolution might select against.

fact, this downplays the severity of our situation. It is not just that we are unreliable when it comes to assessing our own capacity to recognize faces of other races—off the mark in some random fashion. Rather, our self-assessment is systematically distorted: We regard ourselves as better off than we in fact are at recognizing faces of other races. In this way, we are akin to compasses whose arrows are consistently skewed in the wrong direction.⁷¹ In short, by broadening our conception of the CRE, a two-factor approach puts us in a position to appreciate that there is something epistemically defective about it. Blame lies not with the recognitional capacity itself, but with our faulty assumption concerning the way it works.

⁷¹ This analogy comes from Egan and Elga (2005).

References

- Abid, G. (2021). Recognition and the perception-cognition divide. *Mind & Language*.
- Adams, W., Graf, E., & Ernst, M. (2004). Experience can change the “light-from-above” prior. *Nature Neuroscience*, *7*(10), 1057–1058.
- Aggleton, J. (1985). One-trial object recognition by rats. *The Quarterly Journal of Experimental Psychology*, *37*, 279-294.
- Aggleton, J., Vann, S., Denby, C., Dix, S., Mayes, A., Roberts, N., & Yonelinas, A. (2005). Sparing of the familiarity component of recognition memory in a patient with hippocampal pathology. *Neuropsychologia*, *43*, 1810-1823.
- Allen, K. (2013). Blur. *Philosophical Studies*, *162*, 257-273.
- American Psychological Association. (2023). *APA Dictionary of Psychology*.
- Baker, K., Laurence, S., & Mondloch, C. (2017). How does a newly encountered face become familiar? The effect of within-person variability on adults' and children's perception of identity. *Cognition*, *161*, 19-30.
- Beck, J. (2018). Marking the perception–cognition boundary: the criterion of stimulus-dependence. *Australasian Journal of Philosophy*, *96*(2), 319-334.
- Begby, E. (2011). Review of Tyler Burge, *origins of objectivity*. *Notre Dame Philosophical Reviews*, *2*.
- Bengtsson, C., Persson, M., & Willenhag, P. (2005). Gender and overconfidence. *Economics Letters*, *86*, 199-203.
- Bernstein, M., Young, S., & Hugenberg, K. (2007). The cross-category effect: Mere social categorization is sufficient to elicit an own-group bias in face recognition. *Psychological Science*, *18*(8), 706-712.
- Block, N. (1996). How heritability misleads about race. *Boston Review*, *20*(6), 30-35.
- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, *30*, 481-548.
- Block, N. (2014). Seeing-as in the light of vision science. *Philosophy and Phenomenological Research*, *89*(3), 560–572.
- Brady, T., Konkle, T., Alvarez, G., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, *105*(38), 14325-14329.
- Browning, P., Gaffan, P., Crosson, P., & Baxter, M. (2010). Severe scene learning impairment, but intact recognition memory, after cholinergic depletion of inferotemporal cortex followed by fornix transection. *Cerebral Cortex*, *20*, 282-293.
- Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, *77*(3), 305–327.

- Burge, T. (2010). *Origins of objectivity*. New York: Oxford University Press.
- Burgund, E., & Marsolek, C. (2000). Viewpoint-invariant and viewpoint-dependent object recognition in dissociable neural subsystems. *Psychonomic Bulletin & Review*, 7, 480-489.
- Burton, A., Jenkins, R., Hancock, P., & White, D. (2005). Robust representations for face recognition: the power of averages. *Cognitive Psychology*, 51(3), 256-284.
- Burton, A., Jenkins, R., & Schweinberger, S. (2011). Mental representations of familiar faces. *British Journal of Psychology*, 102(4), 943-958.
- Burton, M., Schweinberger, S., Jenkins, R., & Kaufmann, J. (2015). Arguments against a configural processing account of familiar face recognition. *Psychological Science*, 10(4), 482-496.
- Byrne, A. (2005). Perception and conceptual content. In Steup, M., Turri, J., & Sosa, E., (Eds.), *Contemporary debates in epistemology* (pp. 231–250). Oxford: Wiley-Blackwell.
- Campbell, J. (2016). The problem of spatiality for a relational view of experience. *Philosophical Topics*, 44, 105-120.
- Chalmers, D. (2004). The representational character of experience. In B. Leiter (Ed.), *The Future of Philosophy* (pp. 153-181). Oxford: Oxford University Press.
- Chalmers, D. (2011). Verbal disputes. *Philosophical Review*, 120, 515–566.
- Chen, G., Chen, K., Knox, J., Inglis, J., Bernard, A., Martin, S., Justice, A., McConlogue, L., Games, D., Freedman, S., & Morris, R. (2000). A learning deficit related to age and β amyloid plaques in a mouse model of Alzheimer's Disease. *Nature*, 408, 975-979.
- Chiroro, P & Valentine, T. (1995). An investigation of the contact hypothesis of the own-race bias in face recognition. *The Quarterly Journal of Experimental Psychology*, 48(4), 879-894.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181-204.
- Cleary, A. (2008). Recognition memory, familiarity, and déjà vu experiences. *Current Directions in Psychological Science*, 17, 353-357.
- Colibazzi, T. et al. (2010). Neural systems subserving valence and arousal during the experience of induced emotions. *Emotion*, 10, 377-389.
- Coltheart, M. (2010). The neuropsychology of delusions. *Annals of the New York Academy of Sciences*, 1191, 16-26.
- Darby, R. & Caplan, C. (2016). ‘Cat-gras’ delusion: a unique misidentification syndrome and a novel explanation. *Neurocase*, 22, 251-256.
- Darby, R. & Sashank, P. (2016). Lesion-related delusional misidentification syndromes: a comprehensive review of reported cases. *Journal of Neuropsychiatry and Clinical Neurosciences*, 28, 217-222.

- Dodson, C. & Dobolyi, D. (2016). Confidence and eyewitness identifications: The cross-race effect, decision time and accuracy. *Applied Cognitive Psychology*, 30(1), 113–125.
- Dovidio, J. & Gaertner, S. (2000). Aversive racism and selection decisions: 1989 and 1999. *Psychological Science*, 11(4), 315-319.
- Drayson, Z. (2014). The personal/subpersonal distinction. *Philosophy Compass*, 9(5), 338-346.
- Dunning, D. (2011). The Dunning-Kruger effect: On being ignorant of one's own ignorance. *Advances in Experimental Social Psychology*, 44, 247-296.
- Egan, A. & Elga, A. (2005). I can't believe I'm stupid. *Philosophical Perspectives*, 19, 77-93.
- Ellis, H. & Young, A. (1990). Accounting for delusional misidentifications. *British Journal of Psychiatry*, 157, 239-248.
- Ellis, H. et al. (1997). Reduced autonomic responses to faces in Capgras delusion. *Proceedings of the Royal Society B: Biological Sciences*, 264, 1085-1092.
- Ellis, H. & Lewis, M. (2001). Capgras delusion: a window on face recognition. *Trends in Cognitive Science*, 5, 149-156.
- Epstein, R. & Higgins, J. (2007). Differential parahippocampal and retrosplenial involvement in three types of visual scene recognition. *Cerebral Cortex*, 17, 1680-1693.
- Evans, G. (1982). *Varieties of reference*. Oxford: Oxford University Press.
- Fiacconi, C. et al. (2014). Nature and extent of person recognition impairments associated with Capgras syndrome in Lewy body dementia. *Frontiers in Human Neuroscience*, 8, 1-14.
- Firestone, C., & Scholl, B. (2016). Cognition does not affect perception: evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, 39, 1–72.
- Flombaum, J., & Scholl, B. (2006). A temporal same-object advantage in the tunnel effect: facilitated change detection for persisting objects. *Journal of Experimental Psychology: Human Perception and Performance*, 32(4), 840-853.
- Foote, A. & Crystal, J. (2007). Metacognition in the rat. *Current Biology*, 17, 551-555.
- Garcia-Marques, T. et al. (2004). Positivity can cue familiarity. *Personality and Social Psychology Bulletin*, 30, 585-593.
- Gendler, T. (2008). Alief and belief. *The Journal of Philosophy*, 10, 634-663.
- Gennaro, R. (2004). Higher-order theories of consciousness: an overview. In R. Gennaro (Ed.), *Higher-Order Theories of Consciousness: An Anthology* (pp. 1-13). Amsterdam: John Benjamins Publishing Company.
- Goldman, A. (1979). What is justified belief? In Pappas, G. (Ed.), *Justification and belief* (pp. 1-23). Boston: D. Reidel.
- Goodman, R. (2016). Cognitivism, significance and singular thought. *Philosophical Quarterly*, 66, 236-260.

- Goupil, L. & Kouider, S. (2016). Metacognition: behavioral and neural indices of metacognitive sensitivity in preverbal infants. *Currently Biology*, 26, 3038-3045.
- Green, E. J., & Schellenberg, S. (2018). Spatial perception: the perspectival aspect of perception. *Philosophy Compass*, 13(2), 1-16.
- Green, E. J. & Quilty-Dunn, J. (2021). What is an object file? *British Journal for the Philosophy of Science*, 72, 665-699.
- Green, E. J. Forthcoming. The perception-cognition border: a case for architectural division. *The Philosophical Review*.
- Grice, P. (1961). The causal theory of perception. *Proceedings of the Aristotelian Society*, 35, 121-153.
- Harvey-Girard, E., Tweedle, J., Ironstone, J., Cuddy, M., Ellis, W., & Maler, L. (2010). Long term recognition memory of individual conspecifics is associated with telencephalic expression of Egr-1 in the electric fish *Apteronotus leptorhynchus*. *Journal of Comparative Neurology*, 518, 2666-2692.
- Haslanger, S. (2000). Gender and race: (What) are they? (What) do we want them to be? *Noûs*, 34, 31-55.
- Hempel, C. (2002). Two models of scientific explanation. In Y. Balashov (Ed.), *Philosophy of Science: Contemporary Readings* (pp. 45–55). London: Routledge.
- Hirstein, W. & Ramachandran, V. S. (1997). Capgras syndrome: a novel probe for understanding the neural representation of the identity and familiarity of persons. *Proceedings of the Royal Society B: Biological Sciences*, 264, 437-444.
- Hollingworth, A., & Henderson, J. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance*, 28(1), 113-136.
- Hourihan, K., Benjamin, A., & Liu, X. (2012). A cross-race effect in metamemory: predictions of face recognition are more accurate for members of our own race. *Journal of Applied Research in Memory and Cognition*, 1(3), 158–162.
- Huang, G., Ramachandran, S., Lee, T. S., & Olson, C. (2018). Neural correlate of visual familiarity in macaque area V2. *The Journal of Neuroscience*, 38(42), 8967– 8975.
- Hummel, J. (2013). Object recognition. In Reisburg, D. (Ed.), *Oxford handbook of cognitive psychology* (pp. 32–46). Oxford: Oxford University Press.
- Innocence Project. (2008). The role of race in misidentification. Retrieved from: <https://innocenceproject.org/the-role-of-race-in-misidentification/>
- Innocence Project. (2020). How eyewitness misidentification can send innocent people to prison. Retrieved from: <https://innocenceproject.org/how-eyewitness-misidentification-can-send-innocent-people-to-prison/>
- James, W. (1890). *The principles of psychology, volume I*. New York: Holt.

- Jeshion, R. (2010). Singular Thought: acquaintance, semantic instrumentalism, and cognitivism. In R. Jeshion (Ed.), *New Essays on Singular Thought* (pp. 105-40). Oxford: Oxford University Press.
- Kelly, D., Quinn, P., Slater, A., Lee, K., Ge, L., & Pascalis, O. (2007). The other-race effect develops during infancy: evidence of perceptual narrowing. *Psychological Science, 18*(12), 1084-1089.
- Kind, A. (2014). The case against representationalism about moods. In U. Kriegel (Ed.), *Current Controversies in Philosophy of Mind* (pp. 113–134). New York: Routledge Press.
- Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology, 77*(6), 1121–1134.
- Lande, K. (2018). The perspectival character of perception. *The Journal of Philosophy, 115*(4), 187–214.
- Langdon, R. & Coltheart, M. (2000). The cognitive neuropsychology of delusions. *Mind & Language, 15*, 184-218.
- Larzabal, C., Tramoni, E., Muratot, S., Thorpe, S., & Barbeau, E. (2018). Extremely long-term memory and familiarity after 12 years. *Cognition, 170*, 254-262.
- Lawlor, K. (2001). *New thoughts about old things: cognitive policies as the ground of singular concepts*. New York: Garland Publishing.
- Leek, E., & Johnston, S. (2006). A polarity effect in misoriented object recognition: the role of polar features in the computation of orientation-invariant shape representations. *Visual Cognition, 13*(5), 573-600.
- Levin, D. (2000). Race as a visual feature: using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General, 129*(4), 559-574.
- Levy, N. (2015). Neither fish nor fowl: implicit attitudes as patchy endorsements. *Noûs, 49*(4), 800-823.
- Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy, 50*, 249-258.
- Lewis, M. et al. (2001). Autonomic responses to familiar faces without autonomic responses to familiar voices: evidence for voice-specific Capgras delusion. *Cognitive Neuropsychiatry, 6*, 217-228.
- Longmore, C., Liu, C. H., & Young, A. (2008). Learning faces from photographs. *Journal of Experimental Psychology: Human Perception and Performance, 34*(1), 77-100.
- Lucchelli, F. & Spinnler, H. (2008). A reappraisal of person recognition and identification. *Cortex, 44*, 230-237.

- Lupyan, G. (2015). Cognitive penetrability of perception in the age of prediction: predictive systems are penetrable systems. *Review of Philosophy and Psychology*, 6, 547–569.
- Lyons, J. (2005). Clades, Capgras, and perceptual kinds. *Philosophical Topics*, 33, 185-206.
- Mandelbaum, E. (2016). Attitude, inference, association: On the propositional structure of implicit bias. *Noûs*, 50(3), 629-658.
- Malinowska, J. (2016). Cultural neuroscience and the category of race: the case of the other-race effect. *Synthese*, 193, 3865-3887.
- Mandelbaum, E. (2018). Seeing and conceptualizing: modularity and the shallow contents of perception. *Philosophy and Phenomenological Research*, 97(2), 267-283.
- Mandler, G. (1980). Recognizing: the judgement of previous occurrence. *Psychological Review*, 87, 252-271.
- Mandler, G. (2008). Familiarity breeds attempts: a critical review of dual-process theories of recognition. *Perspectives on Psychological Science*, 3(5), 390-399.
- McDowell, J. (1994). *Mind and world*. Cambridge: Harvard University Press.
- McGinn, C. (1989). Can we solve the mind-body problem? *Mind*, 98, 349-366.
- Mehrabian, A. & Russell, J. (1974). *An approach to environmental psychology*. Cambridge: MIT Press.
- Meissner, C. & Brigham, J. (2001). Thirty years of investigating the own-race bias in memory for faces: a meta-analytic review. *Psychology, Public Policy, and Law*, 7(1), 3-35.
- Melcher, D. (2006). Accumulation and persistence of memory for natural scenes. *Journal of Vision*, 6(2), 8-17.
- Michel, C., Rossion, B., Han, J., Chung, C., & Caldara, R. (2006). Holistic processing is finely tuned for faces of one's own race. *Psychological Science*, 17(7), 608-615.
- Michel, M. and Morales, J. (2019). Minority reports: consciousness and the prefrontal cortex. *Mind & Language*, 35, 1-21.
- Millar, A. (2008). Perceptual-recognitional abilities and perceptual knowledge. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, Action, Knowledge* (pp. 330-347). Oxford: Oxford University Press.
- Monin, B. (2003). The warm glow heuristic: when liking leads to familiarity. *Journal of Personality and Social Psychology*, 85, 1035-1048.
- Morales, J., Bax, A., & Firestone, C. (2020). Sustained representation of perspectival shape. *Proceedings of the National Academy of Sciences*, 117(26), 14873-14882.
- Munton, J. (2019). Perceptual skill and social structure. *Philosophy and Phenomenological Research*, 99(1), 131-161.

- Munton, J. (Forthcoming). Bias in a biased system: visual perceptual prejudice. In Ballentyne, N. & Dunning, D. (Eds.), *Bias, reason and enquiry: New perspectives from the crossroads of epistemology and psychology*. Oxford: Oxford University Press.
- Ng, W. & Lindsay, R. (1994). Cross-race facial recognition: failure of the contact hypothesis. *Journal of Cross-Cultural Psychology*, 25(2), 217-232.
- Noë, A. (2002). Is the visual world a grand illusion? *Journal of Consciousness Studies*, 9(5-6), 1-12.
- Noë, A. (2003). Causation and perception: The Puzzle Unraveled. *Analysis*, 63, 93-100.
- Noë, A. (2004). *Action in perception*. Cambridge: MIT Press.
- Orlandi, N. (2014). *The innocent eye: why vision is not a cognitive process*. New York, NY: Oxford University Press.
- Palmero, R., Rossion, B., Rhodes, G., Laguesse, R., Tez, T., Hall, B., Albonico, A., Malaspina, M., Daini, R., Irons, J., Al-Janabi, S., Taylor, L., Rivolta, D., & McKone, E. (2017). Do people have insight into their face recognition abilities? *The Quarterly Journal of Experimental Psychology*, 70(2), 218-233.
- Papesh, M & Goldinger, S. (2010). A multidimensional scaling analysis of own- and cross-race face spaces. *Cognition*, 116(2), 283-288.
- Papineau, D. (2013). Comments on François Recanati's mental files: doubts about indexicality. *Disputatio*, 5, 159-175.
- Peacocke, C. (1983). *Sense and content*. Oxford: Oxford University Press.
- Peacocke, C. (1992). *A study of concepts*. Cambridge: MIT Press.
- Perry, J. (1980). A problem about continued belief. *Pacific Philosophy Quarterly*, 61, 317-332.
- Peters, M. et al. (2017). Does unconscious perception really exist? Continuing the ASSC20 debate. *Neuroscience of Consciousness*, 3, 1-11.
- Phillips, B. (2017). The shifting border between perception and cognition. *Noûs*, 53(2), 316-346.
- Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341-423.
- Quilty-Dunn, J. Forthcoming. Attention and encapsulation. *Mind & Language*.
- Recanati, F. (2012). *Mental files*. Oxford: Oxford University Press.
- Redfern, A., & Benton, C. (2019). Representation of facial identity includes expression variability. *Vision Research*, 157, 123-131.
- Reisenzein, R. (1994). Pleasure-arousal theory and the intensity of emotions. *Journal of Personality and Social Psychology*, 67, 525-539.

- Rhodes, G., Lie, H., Ewing, L., Evangelista, E., & Tanaka, J. (2010). Does perceived race affect discrimination and recognition of ambiguous-race faces? A test of the sociocognitive hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*(1), 217-223.
- Rhodes, M., Sitzman, D., & Rowland, C. (2013). Monitoring and control of learning own-race and other-race faces. *Applied Cognitive Psychology*, *27*(5), 553-563.
- Ricker, T., & Cowan, N. (2010). Loss of visual working memory within seconds: the combined use of refreshable and non-refreshable features. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*(6), 1355-1368.
- Rodin, M. (1987). Who is memorable to whom: a study of cognitive disregard. *Social Cognition*, *5*, 144-165.
- Rosch, E. (1978). Principles of categorization. In Rosch, E., & Lloyd, B. (Eds.), *Cognition and categorization* (pp. 27–48). Oxford: Lawrence Erlbaum.
- Rose, S., Feldman, J., & Jankowski, J. (2004). Infant visual recognition memory. *Developmental Review*, *24*, 74-100.
- Russell, B. (1921). *The analysis of mind*. London: Allen and Unwin.
- Russell, J. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, *110*, 145–172.
- Sangrigoli, S., Pallier, C., Argenti, A., Ventureyra, V., & de Schonen, S. (2005). Reversibility of the other-race effect in face recognition during childhood. *Psychological Science*, *16*(6), 440-444.
- Sarkar, S. (1992). Models of reduction and categories of reductionism. *Synthese*, *91*, 167-194.
- Schellenberg, S. (2008). The situation-dependency of perception. *The Journal of Philosophy*, *105*(2), 55-84.
- Schellenberg, S. (2018). *The unity of perception*. Oxford: Oxford University Press.
- Schroeter, L. (2007). Illusion of transparency. *Australasian Journal of Philosophy*, *85*, 597-618.
- Schwartz, L., & Yovel, G. (2019). Learning faces as concepts rather than percepts improves face recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *45*(10), 1733-1747.
- Schwitzgebel, E. (2006). Do things look flat? *Philosophy and Phenomenological Research*, *72*(3), 589-599.
- Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.
- Simons, D., & Rensink, R. (2005). Change blindness: past, present, and future. *Trends in Cognitive Science*, *9*(1), 16-20.
- Smith, A. D. (2008). Translucent experiences. *Philosophical Studies*, *140*, 197-212.

- Spikman, P., Berg, I., & Deelman, B. (1995). Spared recognition capacity in elderly and closed-head-injury subjects with clinical memory deficits. *Journal of Clinical and Experimental Neuropsychology*, *17*, 29-34.
- Standing, L., Conezio, J., & Haber, R. (1970). Perception and memory for pictures: single-trial learning of 2500 visual stimuli. *Psychonomic Science*, *19*(2), 73-74.
- Strawson, P. F. (1959). *Individuals*. London: Methuen.
- Strawson, P. F. (1974). *Subject and predicate in logic and grammar*. London: Methuen.
- Stroeymeyt, N., Guerrieri, F. J., van Zweden, J. S., & d’Ettorre, P. (2010). Rapid decision making with side-specific perceptual discrimination in ants. *PLoS ONE*, *5*.
- Swinburne, R. (1968). *Space and time*. New York: St. Martin’s Press.
- Tarr, M., & Hayward, W. (2017). The concurrent encoding of viewpoint-invariant and viewpoint dependent information in visual object recognition. *Visual Cognition*, *25*(1-3), 100-121.
- Tullis, J., Benjamin, A., & Liu, X. (2014). Self-pacing study of faces of different races: metacognitive control over study does not eliminate the cross-race recognition effect. *Memory & Cognition*, *42*, 863-875.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge: MIT Press.
- Volpe, B., Holtzman, W., & Hirst, W. (1986). Further characterization of patients with amnesia after cardiac arrest: preserved recognition memory. *Neurology*, *36*, 408-411.
- Wagner, M. (2006). *The geometries of visual space*. Mahwah: Lawrence Erlbaum.
- Whittlesea, B. & Williams, L. (2000). The source of feelings of familiarity: the discrepancy attribution hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 547–565.
- Wittgenstein, L. (1953/2009). *Philosophical investigations*. Translated by G. E. M. Anscombe, P. M. S. Hacker, & J. Schulte. Revised 4th edition by P. M. S. Hacker & J. Schulte. Oxford: Wiley-Blackwell.
- Wittgenstein, L. (1969). *The blue and brown books*. New York: Oxford University Press.
- Wu, W. (2013). Visual spatial constancy and modularity: does intention penetrate vision? *Philosophical Studies*, *165*, 647-669.
- Wu, W. (2014). Against division: consciousness, information and the visual streams. *Mind & Language*, *29*, 383-406.
- Yonelinas, A. (2010). Recollection and familiarity: examining controversial assumptions and new directions. *Hippocampus*, *20*(11), 1178–1194.
- Young, S., Hugenberg, K., Bernstein, M., & Sacco, D. (2012). Perception and motivation in face recognition: A critical review of theories of the cross-race effect. *Personality and Social Psychology Review*, *16*(2), 116-142.

Zahorik, P., & Wightman, F. (2001). Loudness constancy with varying sound source distance. *Nature Neuroscience*, *4*(1), 78-83.

Zajonc, R. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, *9*, 1-27.

Zhang, W., & Luck, S. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, *20*(4), 423-428.

Zhou, X., & Jenkins, R. (2020). Dunning-Kruger effects in face perception. *Cognition*, *203*, 1-12.