CHAPTER SEVEN

PSYCHOLOGICAL CAPACITY AND POSITIVE EPISTEMIC STATUS

Peter Graham[1], University of California, Riverside

More epistemologists than ever recognize and appreciate the plurality of positive epistemic statuses. Knowledge, of course, is the central case. But there are many others. In general, positive epistemic statutes are goods or successes understood in terms of promoting truth and avoiding error. Knowledge is clearly such a good. True belief itself is another. Over the last forty or so years epistemologists have uncovered a number of such statuses—just think of all of the various theories of 'epistemic justification' they've proposed. Nearly every one isolates and at least partially explicates a good or success understood in terms of promoting truth and avoiding error.

An analogy may help. In the first half of the twentieth century philosophers of language widely held that singular terms—definite descriptions, names, indexicals, demonstratives and demonstrative phrases—all worked in the same way; they all worked the way definite descriptions work: 'singular reference' was just one thing. We now recognize that singular reference is many; singular terms don't all work the same way. Likewise there isn't just one kind of positive epistemic status.

Traditionally epistemology investigates the natures, powers and limits of our psychological capacities (Hatfield 1998). Different positive epistemic statuses correlate with different psychological capacities. Some apply, in the first instance, to believers, while

others apply, in the first instance, to beliefs. Some apply to higher non-human animals like apes and maybe dolphins, as well as small children and ordinary human adults. Others apply only to reflective, mature humans. To paraphrase Adam Leite, just as mature humans have abilities animals and young children lack, mature humans enjoy positive epistemic statuses animals and children lack. "Different abilities allow for different statuses" (Leite 2008: 422).

Given the plurality, many epistemologists now self-consciously see 'epistemic justification' as a largely technical term (BonJour 2001: 49; Feldman 2008: 346). They use it as shorthand for a particular positive epistemic status they are interested in isolating and explicating. Many use it functionally for the normativity constitutive of knowledge; on this use, knowledge analytically entails justification. Many externalists tend to use 'justification' this way, in part because of the general recognition that knowledge involves a number of externalist elements or factors (e.g. Alvin Goldman, John Greco, and Ernest Sosa).

Though some internalists also insist that justification is necessary for knowledge internalists tend to connect their use of 'justification' to evidence, good reasons, rational procedures, justifying arguments, or responsible inquiry and reflection (Robert Audi, Carl Ginet, Adam Leite).

On internalist uses it is at least an open question whether knowledge entails justification. Indeed, a number of externalists grant the internalist use of 'justification' and then go on to argue that knowledge does not entail justification (Unger 1968; Dretske 1981; Kornblith 2009). Some internalists do so too. With the advent of reliabilist theories of knowledge, Robert Audi was one of the first to draw a sharp distinction between knowledge and justification (1987; 1988; 2011). On Audi's view, the externalist has the upper hand when explicating knowledge. But when it comes to justification, the internalist wins the day.

Since Audi's and other like-minded proposals in the 1980s, matters have grown more complicated, for two reasons. First, as I've already noted, there's the growing recognition of the plurality of positive epistemic statuses; there isn't just 'externalist knowledge' and 'internalist justification.' Second, for each positive epistemic status, one may naturally query whether it is entirely externalist, entirely internalist, or somewhere in between (BonJour 2001: 51; Kornblith 2010: 8).

A complete epistemology would sort out the whole plurality of positive epistemic statuses relying on a firm understanding of our psychological capacities and kinds, then isolate and explicate each one in turn, laying bare their internalist and externalist elements or dimensions. It would explicate knowledge, rationality, reasonableness, warrant, evidence, justification, virtue, responsibility, reliability, and others, including wisdom and understanding.

I shall not attempt a complete epistemology here; that would take a rather large book if not an entire career. Instead I shall undertake a more modest task. I shall sort out four levels of mind or psychological capacity and then focus on just one distinctive positive epistemic status, paradigmatically enjoyed by perceptual belief. The idea is to isolate a distinctive role for an epistemic status to play, an open spot for a good or success understood in terms of promoting truth to occupy. I shall call it *warrant* (cf. Burge 1993; 2003).

I shall discuss in some detail four levels of mind or psychological capacity. I do this partly for its own sake; getting the psychology right is philosophically worthwhile. But also for epistemology, for different epistemic statuses correlate with different psychological capacities. If one fails to appreciate psychological differences, one is apt to miss, blur, misunderstand or misconstrue various positive epistemic statuses. And it is by appreciating some of the details that one appreciates the differences. And by appreciating the differences one is more likely to appreciate the point of isolating warrant, especially perceptual warrant.

1. FOUR LEVELS OF PSYCHOLOGICAL CAPACITY

Human and non-human animals use all sorts of psychological mechanisms with various degrees of complexity and sophistication for various purposes. Providing a taxonomy of possible psychological mechanisms and their relations would probably require a textbook, and given the state of our knowledge in cognitive and comparative psychology it would surely be incomplete. Regardless, relying on results in the cognitive sciences and some armchair reflection, we can roughly distinguish a handful of psychological mechanisms or capacities. Reflection on cases, experiments, philosophical theory and psychological science all contribute to advancing our understanding of our psychological capacities and correlated epistemic statuses.

In this section I sort out four levels of psychological capacity: functional sensory-registration of information; objective perceptual representation of distal objects, attributes and relations; propositional thinking and reasoning; and critical reasoning (Burge 2003; 2009; 2010).

1.1 Sensory Registration

So there is a sense in which steel 'carries information' about the presence of oxygen in the environment: the steel wouldn't rust unless it was in contact with oxygen (Dretske 1981). Rusting obviously isn't a psychological process. Psychology doesn't begin with information carrying. Psychology begins with information used to guide biologically useful behavior.

Changes on the amoeba correspond to changes in heat, light, and acidity. The surface of the amoeba carries information about features of its environment. The surface, however, doesn't just carry information. Unlike steel, the amoeba makes use of the information carried. For unlike steel, organisms behave. Change the acidity of the water and the amoeba will change direction. Heat up one side or shine a bright light and it will turn away. The movements of the amoeba have obvious functional benefits. Heat and acidic water isn't any good for the amoeba. To fulfill its biological needs the amoeba needs to move in various ways in response to changes in its environment. The amoeba thus recruits information carried about its environment by changes on its surface—its sensory receptors—to guide its behavior in fulfilling its biological needs. To sense isn't just to carry information, but to make use of information carried by proximal stimulation in the control of biologically useful behavior.

In general a sensory system carries or 'registers' information about the whole organism's environment that assists the organism in fulfilling its whole organism biological needs. Sensory systems are systems whose function is to carry information about the environment that in turn controls or contributes to biologically beneficial behavior of the whole organism. A sensation or sensory registration is then a state of such a system; it's a state that's supposed to carry specific information about the organism's environment (Burge 2010; Dretske 1986; Godfrey-Smith 1996). Proximal information-carrying stimuli produce biologically useful behavior.

Many philosophers connect sensation with conscious sensory experience. From the first-person perspective, this can seem rather

obvious, for only conscious sensations are accessible to first-person reflection. But is phenomenal consciousness present whenever sensory registration occurs? Is there something it is like for the amoeba to sense the presence of light or a physical obstacle that impedes its movement? I have my doubts. The science of consciousness does not, to my knowledge, have an answer. Sensory registration may occur long before the onset of phenomenal consciousness. Though phenomenal consciousness often accompanies human sensory registrations, it often occurs, in particular cases, without phenomenal consciousness. The connection between functional sensory information carrying and phenomenal consciousness is, to my mind, unclear. Sensory *registrations* are not always sensory *experiences*.

Sensory systems are ubiquitous in the animal kingdom; you're not a member of the kingdom unless you sense your environment, in one way or another. They are present in all species of mammals, reptiles, amphibians, birds, fish, insects, and arachnids. Having a sensory system is one way an organism fulfills its biological needs. And for many organisms sensory systems are enough; there is no need for further capacities or faculties to functionally adapt to its environment.

1.2 Perceptual Representation

Many animals also perceive distal objects, attributes, and relations; they perceptually represent their world. The amoeba only *senses* the warmth of the water. The eagle *sees* its prey emerging from the water.

What's the difference between a sensory system and a perceptual system? What else is involved in perceiving the environment beyond sensing it? What's the psychological mechanism perceptual systems deploy that sensory systems lack?

One crucial element involves *perceptual constancies* (Walsh and Kulikowski 1998; Burge 2003; 2010). Imagine a box moving towards you. As it gets closer, the two-dimensional image on your retina—the proximal stimuli on your sensory organs—will continually grow: a series of differing proximal stimuli follow one after another. But nonetheless the three-dimensional size of the box continually looks exactly the same size to you. You continue to perceptually represent the distal object—the box—as having the same attribute—the same size—despite continual changes in proximal stimulation. This is an example of what's

called size constancy. Your perceptual representation is *constant*—it keeps representing the box as having the same size—despite changes in the sizes of the images on your retina. And for every type of thing we perceive—objects, size, shape, motion, location, distance, color—we rely on constancy mechanisms, mechanisms that produce representations of distal objects, attributes and relations as being the same despite various changes in sensory, proximal stimuli. It also works the other way. Two identical sensory stimuli in different environmental conditions can produce differing perceptual representations. Perceptual systems thus involve transitions from sensory stimuli that underdetermine the features of the distal stimuli to perceptual representations of the distal stimuli. Understanding these transitions is in large part what perceptual psychology is all about. Perceptual psychology asks how we form representations of a stable world from ever changing sensory input.

Sense perceptual systems involve constancy mechanisms, but mere sensory systems do not. Hence the idea: perception differs from sensory registration in 'distalizing,' in transitioning from proximal stimuli to representations of distal objects, attributes, and relations.

Sensory systems and sense-perceptual systems are two psychological mechanisms with obvious biological benefits. Many organisms only have sensory systems. Many have both. We certainly do. Fish and frogs do too. Bees do. Fleas? Worms? I don't know. Amoeba? Definitely not.

Perceptual systems produce perceptual representations. These are representations of particulars: of particular objects, particular attributes, or particular relations. But they also group or categorize the particulars as instances of groups or categories. That object is represented as *that brown body*. That physical object is represented as *that cube*, or as *that moving round body*. Perceptual representations are accurate when the particular represented has the attribute or relation the representation represents it as having (Burge 2003; 2009; 2010). Many philosophers connect perceptual representation with phenomenally conscious perceptual experience. From the first-person perspective, this can seem rather obvious, for only conscious perceptual representations are accessible to first-person reflection. But is phenomenal consciousness present whenever perceptual representation occurs?

Is phenomenal consciousness present whenever perceptual representation occurs? Is there something it is like for the bee to perceive the direction of food? I have my doubts. The science of consciousness does not, to my knowledge, have an answer. Perceptual representation

may occur long before the onset of phenomenal consciousness. Though phenomenal consciousness often accompanies human perceptual representation, it sometimes occurs, in particular cases, without phenomenal consciousness. Just think of blindsight. The connection between sense perception and phenomenal consciousness is, to my mind, unclear. Not every perceptual representation is conscious, and probably not every perceptual representation that is conscious is necessarily accessible in mature humans to reflection. Perceptual *representations* are not always perceptual *experiences*.

There is a rather large literature on the form and content of perceptual representations. Are they analog or digital? Do they represent like pictures, maps, or images, or more like sentences? Do they have non-conceptual content, or is all content conceptual? Are their contents abstract Fregean *Sinn*, sets of possible worlds, or structured Russellian propositions?

Burge (2009; 2010) argues that perceptual representations have non-propositional structures or forms. Their forms are more like singular noun phrases—e.g. *that round ball, that moving blue cube*—as opposed to singular subject-predicate sentences—e.g. *that is a round ball*, or *that is an object and is moving and is blue and is a cube*. And so, in a sense, a perceptual representation with one kind of structure may have the same 'worldly' veridicality condition or the same 'content' as a sentence with a different kind of structure, e.g. *that round body* is accurate just in case the referent is round and is a body, and *that is a round ball* is accurate just in case the referent is round and is a ball. Though, in another sense, because they have different structures they have different contents, and so different veridicality conditions.

Having this kind of structure or form means that perceptual representations are non-propositional; they have non-propositional structure. Propositional structures are like the structures of indicative sentences. A representation with a structure like a singular noun phrase doesn't have that kind of structure. Propositional thoughts and other propositional attitudes, on the other hand, paradigmatically have propositional structures. Perceptual *representations* thus differ from perceptual *beliefs* in having different structures. I accept this view.

1.3 Propositional Thinking

What is propositional structure? It's not just reference and predication, for in a sense that's present in perceptual representations.

They refer to or single out a particular and group or categorize the particular as having various attributes.[2]

One difference between representations with propositional structure involves the ability to reason or draw logical inferences. Logical constants like OR and NOT operate on representations with propositional structure; they are basic connectives in truth-functional, propositional logic. If you attribute the psychological capacity to reason or draw logical inferences to a creature, then you are attributing to the creature the capacity to reason using logical constants, and so attributing to the creature representations with propositional structure. So one contrast between perceptual representations and propositionally structured representations involves the capacity to reason, to draw logical inferences, and so to represent logical relations.

Capacities to reason may come in degrees or in various packages. That is, one may only have the capacity to draw simple propositional inferences like disjunctive syllogism and simple forms of conditional reasoning, without the capacity to represent generality in the form of existential and universal quantification. The capacity may be domain specific. One creature's inferential capacity may have severe limitations. Another creature may enjoy an unlimited, universal inferential competence, limited only by time and energy.

If we use 'reasons' rather narrowly to denote representations with propositional structure—so that only reasons are involved in propositional reasoning—then it follows that the transition from a non-propositional perceptual representation to a propositional perceptual belief is not reasoning. On this usage, only propositional thinking constitutes reasoning.

Is all propositional thinking phenomenally conscious? No. A good deal is fast and automatic, beneath the level of consciousness. Some is slow and ponderous but still below consciousness. What about conscious thinking? Is it always phenomenally conscious? I don't know. But I do know the issue is currently disputed.

There is a lively debate with a very long history over the rational (logical reasoning) capacities of non-human animals (Hurley and Nudds 2006; Rescorla 2009). In the contemporary debate, it is universally agreed that animals employ both learned and innate strategies for getting around their environments. A principle capacity involves the capacity to associate stimuli. Ring a bell often enough in the presence of food, and the animal will associate a perceptual

representation (or perhaps, for the behaviorist, just an auditory stimuli) of the bell with an anticipatory representation of food (or, again, just a salivating response). Associating is one 'intelligent' capacity for learning about and navigating one's environment. And it is widely thought that no matter how much representation actually goes on in animals when associating, associating isn't reasoning or relying on inference. Logical reasoning is a different kind of psychological capacity. So the contemporary debate concerns, in large part, the reasoning or inferential abilities of non-human animals. Do they, or don't they, reason?

There is some experimental evidence that suggests that apes and chimpanzees reason. In front of chimpanzees, Premack and Premack (1994) took two boxes and placed an apple in one and a banana in another. Later the chimpanzee would witness the experimenter eating the banana. Then when given the opportunity to go pick a fruit from the two boxes, the chimps would go right to the box with the apple. They concluded that the chimps reasoned something like this: there is an apple in box *A* and a banana in box *B*. But there is no longer a banana in *B*, so there's just an apple in *A*. That's why they went right for *A*. Animals that don't reason like this, but presented with the same information, might still look for a banana in box *B*, or might only slowly make their way to box *A*.

Josep Call (2006) set out to extend this research with experiments involving apes. But this time he wanted to compare and rule out the rival associative hypothesis. So he put two opaque cups in front of an ape, one with food and the other without. He then shook them. One made noise because of the food; the other did not. Shake the one with food, and the apes tended to pick that one. Shake the one without food and the apes tended to pick the other one. Call thinks they reasoned causally. In the first case, 'it's the food making the noise, so grab the shaking cup'. In the second 'when there's no noise, there's no food in the shaking cup, so grab the other one'.

Call ruled out the rival associative hypothesis by again placing two cups in front of the ape, one with food and one without. But then he conditioned a sound stimulus with the cup with food. Perhaps he tapped the cup with the food instead of shaking it, or played a recorded sound. It turns out that the apes performed worse in the 'associative' condition; they didn't reliably tend to pick the cup with food, despite the conditioning. They were worse at associating the stimuli than reasoning from the cause-effect relationship between the

food, the shaking, and the sound. According to Call, the "apes do not simply associate a sound with the presence of food, but attribute the sound to the food itself, they understand that the food is the cause of the noise" (2006: 222). Apes, he argues, are able to "reason about their physical world" (2006: 220).

Dogs, despite their obvious intelligence, fail these tests. Though they can use sound to detect the presence of food, when you shake the empty cup before them they don't move to the unshaken cup where the food is hidden.

Additionally Watson *et al.* (2001) found that dogs (at least for some tasks) seem to deploy associative search strategies and do not rely on reasoning, where children to the opposite. They predicted that if you trained a dog to search for a ball behind three different screens, then later set the dog off to search behind the screens in the test situation when a ball was not hidden behind any three, the dog would slow down its search. The time taken to move between 1 and 2 would be less than the time taken between 2 and 3; the dogs would slow down as they searched behind the screens. Why? Because if they were relying on association, each time they failed to find a ball the experience would be an extinction trial. 'No ball here? Well, then less likely a ball there'. They then predicted that if you trained children between four and six years of age to perform the same task, the time taken to move between 1 and 2 would be more than the time taken between 2 and 3; the children would speed up their search. Why? Because if the children are reasoning logically then they are reasoning like this: 'there is a ball behind 1, 2, or 3. If there is no ball behind 1, then more likely there is one behind 2 or 3. And if not behind 2, then definitely 3'. And their predictions turned out to be right. They concluded that "dogs rely on associative guidance and children rely, to some degree, on logical guidance when searching for objects that have recently disappeared" (2001: 225). Though very smart, dogs may lack the capacity to reason. Animal intelligence doesn't require the capacity to draw logical inferences.

Non-human animals have a number of surprising cognitive abilities. But only some reason—engage in propositional thinking and propositional inference. When they do, they enjoy propositional thoughts, representations with propositional structure. Paradigmatically beliefs have propositional structure. Though many animals form perceptual (and other kinds of) representations, only some animals reason and form propositional beliefs. And since epistemology (traditionally

conceived) only begins with beliefs, epistemology only begins with animals that reason.

Our first-order representational perspective on the world is largely made up of current and stored perceptual representations and propositional beliefs. Some of these states are phenomenally conscious and some are not. Some are accessible in humans to conscious reflection; some are not. Our first-order representational perspectives are formed and sustained by sub-personal and person-level systems that transition from sensory and perceptual representations to other representational states.

1.4 Critical Reasoning

Many primates reason. We reason all the time. We often reason without much awareness of what we are doing, of what our reasons are for our conclusions, and whether our reasons are any good. We often reason blind (Burge 1996: 99). Blind reasoning isn't blank or empty. It is reasoning that isn't critical reasoning. What's the difference?

Critical reasoning involves an appreciation of reasons *as reasons* (Burge 1996: 99–101). When we critically reason, we weigh and appreciate reasons for or against a proposition or a course of action. When you wonder whether something is true, or whether something is a good idea, you are reasoning critically. When you tell me why something is so, or ask me why I believe what I do, you are reasoning critically. As you read this paper and nod attentively or shake your head in disagreement, you are reasoning critically.

Critical reasoning involves an appreciation of good and bad reasons. We have an appreciation of what follows from what, when something is a good or bad reasons for a conclusion or a course of action. We have a sense of validity, a sense of good or rational support. We have a sense of when something is good evidence for something else.

When reasoning we are normally reasoning about the world, about a first-order subject matter. I'm reasoning about the library's hours and whether I should call first or just set out. But we also think about our reasons as reasons in critical reasoning, and so we think about our propositional attitudes. We assess them as true or false, as well supported or unreasonable. And so crucial to critical reasoning is the capacity to think second-order thoughts, thoughts about thoughts as reasons, as true or false, as well-evidenced or not, both thoughts about our own thoughts and thoughts about the thoughts

of others. In critical reasoning we are aware of reasons as reasons, and relations of reasonable or rational support among reasons.

Critical reasoning involves conscious states and events of the agent. Critical reasoning can only engage accessible states and events. Critical reasoning reviews states currently conscious, such as a current conscious perceptual representation (is the ball really that color?) or current conscious thoughts (I thought Reno was east of Los Angeles, but maybe I am wrong). It might also bring to consciousness stored states and events, or current states and events that are currently just below consciousness. But it cannot review states and events that cannot be made conscious. Much cognitive activity is sub-personal and inaccessible to conscious review. Phenomena like blindsight suggest that even some current perceptual representations—non-sub-personal outputs of the perceptual system—are beyond the purview of critical reason.

When we critically reason we often reach evaluations of our reasons that reinforce our reasons, beliefs, and decisions, but that also lead to changes or revisions in our reasons, beliefs, and decisions. Critical reasoning often leads to changes in our propositional attitudes. It's one way we change our minds. It is how we regulate, revise, and review our perspective or point of view.

Critical reasoning thus involves psychological capacities not present in many higher non-human animals, as well as children, especially three-year-olds and younger. Very young children cannot represent propositional thoughts as thoughts; they cannot represent reasons as reasons; they cannot evaluate reasons as good reasons. They cannot critically reason. When they reason, they reason 'blind.'

Critical reasoning as such only requires the capacity to appreciate reasons as reasons and the ability to criticize reasons and reasoning as reasons and reasoning. Reflective critical reasoning involves more. Reflective critical reasoning "makes use of all the main concepts necessary to a full understanding of essential or fundamental elements in reasoning" (Burge 1996: 98, n. 4). We approximate reflective critical reasoning as we mature and also become more reflective about our reasoning. There is a whole range of critical reasoning between the degree exercised by children five to six years old on the one hand, and fully mature humans who have taken the time to develop their skills at critical reasoning on the other.

A good amount of our critical reasoning may be considerably less than ideal. We're subject to all sorts of biases, and we make any

number of regular mistakes. Critical reasoning courses, especially those informed by current research in social psychology, aim at improving our skills at critical reasoning.

To review: steel rusts but does not sense; amoeba sense but do not perceive; bees perceive but do not reason; some primates and other animals reason but do not critically reason; mature humans sense, perceive, reason, and critically reason.

I have gone on at some length sorting out these four levels of psychological capacity in the conviction that different epistemic statuses correlate with different psychological capacities. If one fails to appreciate psychological differences, one is apt to miss, blur, or flatly misunderstand or misconstrue various positive epistemic statuses. And it is by appreciating some of the details that one appreciates the differences. And by appreciating the differences you'll appreciate the point of isolating warrant, especially perceptual warrant.

2. PERCEPTUAL WARRANT

Epistemologists have used 'epistemic justification' to denote a large number of distinct, though often related, positive epistemic statuses. It's been used to denote responsible belief, belief that results from (or is tied to) the fulfillment of epistemic duties, belief based on the exercise of intellectual virtues, reliable belief, belief based on or supported by evidence, and the normativity constitutive of knowledge. In this section I isolate a spot for the distinctive epistemic status I call warrant, and then briefly sketch my account of perceptual warrant. Warrant arises with perception and the capacity to reason. But first I want to say a few words about another epistemic status that requires the capacity to critically reason.

Call *being justified* the status you achieve when you are in a position to justify your belief. Justifying is an activity that involves the ability to adduce grounds, evidence and reasons as grounds, evidence, and reasons, and so requires the capacity to critically reason. Animals and small children can't do that. And so only those in a position to justify a belief are justified in believing it.

What occurs when we justify a belief? At least three things. First we cite some property or feature of the belief. We may cite accompanying sensory experiences: it feels hot, that's why I think we left the heater on. We may cite memories: I recall meeting Richard last week; that is why I think I know him. We may cite perceptual experiences:

look, it looks as green to me as a green apple; that's why I am sure his new car is green, not blue. We may cite various facts: the earthquake disrupted oil production; that's why prices are about to go up. We are aware of what we cite; we can't cite something we aren't aware of. Just try. As soon as you do cite it, you're aware of it. And if you're not aware of it, you can't cite it. Being able to justify requires conscious access to justifying states, events, and their justifying properties.

Second we appeal to some norm or standard. What kind? One that has to do with truth. When we give a justification, necessarily we aim at citing facts or reasons that provide good grounds or reasons for thinking the belief we are justifying is true. Since we aim at truth when justifying, justifications (*qua* arguments given when justifying) aim at truth. So when we cite a property of the belief in justifying it, and at least implicitly appeal to a standard that the belief falls under, we are implicitly if not explicitly claiming that beliefs that meet that standard are likely to be true. If we didn't think meeting that standard somehow counted towards truth, we wouldn't cite the fact that the belief meets that standard when justifying it. The process of justifying necessarily has truth as its aim (Audi 1988).

And just to be clear, these two are connected; we are aware or claim that the belief meets that standard in virtue of possessing the property cited. Justifying then cites three things: the property of the belief, the norm or standard, and falling under or conforming to the standard.

Ordinarily this is done in language in conversation with others. But we might justify a belief to ourselves, perhaps even non-linguistically. Regardless, justifying exercises our capacity to critically reason, often reflectively so. We do it a good amount of the time. Animals and small children never do, because they can't. Being in a position to justify, and so being justified, only emerges with higher-order, reflective capacities.

Being justified contrasts with being warranted, or warrant. Being justified is a positive epistemic status that correlates with critical reason. Warrant arises prior. Why should we assume, in general, that there exists such a positive epistemic status?

If higher non-human animals and small children have beliefs, then they have true and false beliefs. So trivially their true beliefs enjoy a positive epistemic status: *truth*. And surely also they know things. And so their beliefs enjoy at least two positive epistemic statuses distinct from the property of being justified. Are these the only

two? Or is there another positive epistemic status their beliefs might enjoy?

Yes. Epistemologists in general are willing to attribute to animals and small children a positive epistemic status. Focusing on perceptual belief, one may ask whether relying on a perceptual representation is an epistemically proper response, or whether there's nothing epistemically special about relying on perceptual representations when forming perceptual belief. Put that way, it clearly seems there is. To paraphrase Richard Feldman, we can assess beliefs in terms of whether they are proper responses to the information the individual had to go in when forming the belief. A belief that responds properly to the information is an 'epistemically proper' belief. When it comes to perceptual belief, relying on a perceptual representation is the 'proper response.' Responding in the proper way confers a positive epistemic status. And this is so even if the belief formed falls short of knowledge, and even if it falls short of truth. And since children and animals rely on perception, their beliefs enjoy such a status.

Feldman recognizes that different epistemologists will offer different accounts of what constitutes the proper response. Still they agree on the general idea. "Reliabilists, proper functionalists, evidentialists of various stripes, and others, all agree that there is some notion of a proper response to information (or evidence or stimuli), and that paradigmatic epistemic evaluations are about this. A belief is favorably evaluated when it is a proper response and unfavorably evaluated when it is an improper response" (Feldman 2008: 347).

As I use 'warrant,' a warranted perceptual belief is an 'epistemically' proper response to a perceptual representation, a response that does not entail truth and may fall short of knowledge. Epistemologists may happily agree that perceptual beliefs enjoy warrant, while disagreeing over its constitutively necessary, or constitutively sufficient, or constitutively necessary and sufficient, conditions.

So I think once we have a rudimentary grasp on the nature of perception and perceptual belief, we're all willing to recognize that perceptual beliefs have an epistemology, and so enjoy various positive epistemic statuses, but not just knowledge and true belief. Warrant is an intermediate status; a warranted belief is a 'proper' response to a perceptual representation.

I shall therefore assume that perceptual beliefs paradigmatically enjoy warrants, and that perceptual beliefs are warranted, in part, by perceptual representations. Perceptual warrant is empirical warrant; empirical warrant essentially depends on sense-perceptual representations for its warranting force. Perceptual warrant depends upon the exercise of perception and essentially relies on perceptual representation. Perceptual warrant partly consists in being a proper response to a perceptual representation.

Perceptual belief involves a transition from non-propositional perceptual representation to propositional perceptual belief. Ordinarily the perceptual representation plays a constitutive role in the individuation of the perceptual belief. Perceptual warrant thus partly turns on the transition from the perceptual representation to the perceptual belief. Insofar as we are warranted when relying on perception, both the perceptual representation and the transition from representation to belief contribute to the warrant for the perceptual belief.

I have also assumed that perceptual representations lack propositional structure and so are not reasons, and so not warranting reasons. The transition from a perceptual representation to a perceptual belief is then not an instance of good or warranting reasoning. I know that a number of philosophers tend to use 'reason' more broadly so that non-propositional perceptual representations count as reasons for perceptual beliefs, or that even non-representational facts or worldly states of affairs count as reasons. The word does not matter. What matters is the distinction between representational structures and capacities. Perceptual representations are non-propositional and do not require the capacity to reason. Perceptual beliefs are propositionally structured representations and require the capacity to reason.

Following Burge's terminology, I call non-propositional warrants and contributions to warrant *entitlements*. Perceptual representations entitle perceptual beliefs; they contribute to the overall warrant for a perceptual belief. The transition from a representation to a perceptual belief also plays a role. We are entitled to rely on the normal transition from perception to belief. The normal transition thus contributes to the overall warrant for a perceptual belief; not just any old transition will do. Since the perceptual representation and the transition from representation to belief are both non-propositional, they are both entitlements: non-propositional contributions to warrant.[3]

So far this is mostly terminology. Epistemologists of different persuasions should find nothing objectionable in the terminology.

This is especially true for epistemologists that recognize the existence of positive epistemic statuses that apply in the first instance to beliefs, especially the perceptual beliefs of animals and small children.

But now we can see the point of the terminology. Being justified or justified belief only arises with critical reason. Warrant arises prior, with perception and the capacity to reason. But not all warrant arises from reasons—other warranted beliefs. Entitlements are non-propositional warrants, warrants that are not themselves other warranted beliefs.

There are two contrasting tendencies in most theoretical explications of warrant. One is associated with reliabilism. Reliability theories of warrant emphasize *mind-world* connections, *vertical* connections between the outputs of the belief-forming process or system, on the one hand, and the stretch of the world that comprises the subject matter of the process or system, on the other (Burge 2010: 50–1). Warrant, on this view, entails reliability; most warranted beliefs, necessarily, are true. Warrant actually results in truth, at least for the most part. This first tendency sees warrant as a good route to truth. Like substantive views of justice, good outcomes constitute good procedures.

The second tendency emphasizes good procedure or correct processing in belief-formation; it emphasizes *mind–mind* connections, *horizontal* connections between input representations to the belief-forming system or procedure and the resulting outputs. Rational procedures in logic, mathematics and experimental science provide the model. John Pollock and Joseph Cruz have championed this tendency (Pollock and Cruz 1999; Pollock 1999). Applying the models to perceptual belief formation, the perceptual beliefs of animals and small children would enjoy warrant for perception is a good or correct epistemic procedure. Warrant, at best, properly aims at truth, even if it doesn't actually result in true beliefs. Like procedural views of justice, good procedures stand on their own two feet.

Though contrasting, these two tendencies are not necessarily opposing. I see the value in understanding warrant as first and foremost the result of conforming to correct procedures. However, I agree with the first tendency that warrant is, fundamentally, constitutively associated with reliably getting things right; that's where the connection between truth and warrant lies.

Robert Audi drew a similar contrast (1988). *Ontological* conceptions of 'justification' tied it to reliably connecting with truth. *Teleological* conceptions, on the other hand, tie 'justification' to properly aiming at truth. The first emphasizes outcomes, the second emphasizes goals.

In the remainder of the paper I shall sketch my own account. It explicates perceptual warrant in terms of reliably getting things right in normal conditions. And so in a sense I've got a reliabilist view. But it will also emerge that my account is rather friendly to second tendency views.

This is not the place to argue in detail for my account of perceptual warrant, or to argue against different accounts. It's a big project, only partially completed (Graham 2010a; 2010b; 2011a; 2011b; in preparation-a and -b). But sketching the account should provide an existence proof for the role. A simple de facto reliability theory might have done the same thing, but in a way less plausible, at least to my mind, in part because it would fail to capture the horizontal tendency in explicating warrant, especially perceptual warrant.

My account of perceptual warrant exploits four main ideas: the connection between functions and norms; the connection between goods or successes and the fulfillment of functions or the meeting of norms.

Perceptual warrant is just one kind of warrant. Different kinds may require different accounts. Recall the methodology: I have isolated a role for a positive epistemic status by distinguishing perception and reason, on the one hand, and critical reasoning, on the other. So there is a role for a positive epistemic status that applies, in the first instance, to the perceptual beliefs of animals and small children, a status that falls short of knowledge and like many other statuses does not entail truth, despite being understood in terms of promoting truth and avoiding error. I have also assumed that perceptual beliefs enjoy this status. So we are looking for a good or success understood in terms of promoting truth and avoiding error that perceptual beliefs clearly enjoy, a property of perceptual belief that fills that role.

My account of perceptual warrant and entitlement exploits four main ideas: the connection between functions and norms; the connection between goods or successes and the fulfillment of functions or the meeting of norms; the etiological theory of functions associated with Larry Wright (1973) and Ruth Millikan (1984), among many others; and the empirical claim that human

perceptual systems have, as one of their functions, the function to reliably induce true beliefs.

The first idea is straightforward. Functions are norms or standards for what a functional item is supposed to do (what effect it is supposed to produce). Your car is supposed to take you to work. Your stapler is supposed to staple papers together. Your heart is supposed to pump blood. When your car gets you to work it has fulfilled its function. When you heart pumps blood it does what it is supposed to do.

Cars, hearts and staplers are all also supposed to work or operate a particular way. Functioning (operating, working) normally is another norm or standard your car, heart and stapler may fulfill. Here are two examples that illustrate the difference between function fulfillment and normal functioning. Suppose you take your car into the shop for inspection. Your mechanic may put your car up the lift and then run it at maximum speed. Up the lift your car may be functioning normally (it's in perfect shape) but since it is not taking you anywhere, it is not fulfilling its function. In order to do that it needs to be in normal conditions. You'll have to take it down the lift and drive off the lot to use it to take you where you want to go. Here's the second example. Suppose a surgeon removes your heart from your chest in a complicated, futuristic surgery, and then places it in a vat where it's chemically and electrically stimulated so that it continues to beat and function normally. Now your heart is, as it were, up the lift. But no blood is flowing through; maybe it's orange juice instead, or nothing at all. Your heart is functioning normally but not fulfilling its function. So there are two norms or standards for functional items: function fulfillment and normal functioning.

That's the first idea. The second asserts that fulfillments of functions or the meeting of norms are goods or successes for functional items. It's a good for the heart to pump blood. It's a good or achievement for a heart to function normally. These goods are not necessarily goods for the whole organism, though biological functions are constitutively associated with whole organism goods. Functional goods are not necessarily moral or aesthetic goods either. From the moral point of view, it might be a bad thing for a morally bad creature's heart to function normally. Architecture might fulfill its functions, despite any political and aesthetic consequences.

The third idea is the etiological theory of functions and normal functioning. On the etiological theory, the function of an item is its selected effect. Crudely, the function of a type T is F just in

case F is a selected effect of ancestors of T, where descendants exist because ancestors did F. Think of your eyes. They help you see. And why do we have eyes? Because precursors of our eyes helped our ancestors see. Seeing helped them fulfill their biological needs so that they might produce offspring and thereby produce descendant eyes. The selected effect of ancestors of the type explains why the functional item type exists. Natural selection is one kind of selection. Much learning involves another kind. The item can be a trait like an organ but also a form of behavior or an acquired skill.

The distinction between a function of an item and an 'accidental' side effect partly motivated the etiological account. It's the function of the heart to pump blood, but it also makes regular thumping noises. We have hearts now because they pump blood, not because they make regular noises. Pumping blood is the selected effect, not making noise.

When an item has an etiological function there is a historical explanation for why the item has that function (Millikan 1984). Look at why ancestors were selected for a particular effect. The explanation will include how the item operated in its circumstances so as to produce the selected effect. The way the item worked or operated according to the explanation then counts as *normal functioning*. The circumstances in which it was selected for producing that effect while working that way then count as *normal circumstances*. When an item has an etiological function it is then *a priori* that *ceteris paribus* it will produce its selected effect (it will fulfill its function) often enough in normal conditions when functioning normally. If it didn't produce the effect often enough in normal conditions when functioning normally then it would not have been selected for that effect (selection can't work on effects that don't occur), and then it wouldn't have that effect as a function.

An important consequence follows. Normal functioning (working or operating the way the item is supposed to work or operate) is then constitutively associated with function fulfillment. If it didn't produce the functional effect while working normally, then it wouldn't have that function, and nothing would count as normal functioning. Normal functioning is then partially understood and individuated in terms of function fulfillment; normal functioning *just is* operating the way it operated in normal conditions when fulfilling its function. Normal functioning thus 'encodes' function fulfillment; the latter

types the former. What makes the transition within the trait, system or behavior a normal or correct transition involves producing functional outcomes in normal conditions. 'Horizontal' correctness within the trait or system requires 'vertical' correctness, often enough, in normal conditions; those inner transitions wouldn't be 'horizontally' correct unless they resulted in 'vertically' correct outcomes.

Even so, normal functioning and function fulfillment are token distinct, for one can function normally without function fulfillment. Don't forget the car up the lift. The car functions normally but it doesn't take you where you want to go.

Putting the first three ideas together, an item with a function can meet or satisfy two standards: it can fulfill its function and it can function normally. And when an item has an etiological function, normal functioning is constitutively associated with function fulfillment. And so when an item has an etiological function, both norms that it may meet are constitutively associated with function fulfillment. That's obvious when it fulfills its function; function fulfillment is trivially associated with function fulfillment. It requires understanding the constitutive association between normal functioning and function fulfillment to see why this is so for normal functioning. Since normal functioning 'encodes' function fulfillment, normal functioning is a good or success understood in terms of function fulfillment. Horizontal correctness 'encodes' vertical connections; horizontal correctness is a good or success understood in terms of vertical connections.

So here we have a general theory of goods constitutively associated with functions and norms, especially norms for items with etiological functions. My account of perceptual epistemic entitlement emerges when this general theory is applied to our perceptual belief-forming capacities.

Perceptual beliefs result from transitions from perceptual representations. Perceptual representations are formed through the operation of our sense-perceptual organs and the workings of the perceptual system. Suppose empirically that an etiological function of the primate perceptual system is to form veridical representations reliably. Suppose empirically that an etiological function of the perceptual system in humans is to form true perceptual beliefs reliably. Suppose it does the latter, in part, by doing the former. Then when the human perceptual system is functioning normally in normal conditions it will, *ceteris paribus*, reliably induce true

perceptual beliefs. Normal functioning for the system is then constitutively associated with reliably forming true beliefs; our normally functioning perceptual belief-forming process 'encodes' reliably getting things right. The mind-mind procedure 'encodes' mind-world relations.

Two goods constitutively associated with reliably promoting true belief then apply to our perceptual belief-forming systems. One arises from normal functioning, the other from function fulfillment. In previous work I emphasized the good that arises from normal functioning. I have concluded that *prima facie pro tanto* perceptual epistemic entitlement consists in the normal functioning of the perceptual system and the perceptual belief-forming process, for the overall process has forming true beliefs reliably as an etiological function. Normal functioning is then a good or success understood in terms of promoting truth and avoiding error—a paradigm case of a positive epistemic status—that does not involve reasoning, and so counts as an entitlement, a warrant supporting the resulting perceptual belief (Graham 2010a; 2011a). But for the same reason fulfillment is also a positive epistemic status for the perceptual belief-forming system. One kind or dimension consists in normal functioning, and the other consists in function fulfillment.

And so I have advanced a 'proper function reliabilist' account of perceptual entitlement. Other accounts in the ballpark include Plantinga's 'intelligent design' reliabilist account (Plantinga 1993), and Tyler Burge's account that relies on his distinctive anti-individualist account of the natures of perceptual states (Burge 2003). I discuss Burge's account elsewhere (Graham in preparation-b).

Plantinga's use of 'warrant' suggests that he is really talking about knowledge, and not about the positive epistemic status I have used 'warrant' to isolate. But in fact this is not entirely so. Though Plantinga is not explicit about this, his account involves three grades of warrant. The first requires that the belief-forming process be designed to promote true beliefs in normal conditions, and that the plan be a good one, so that if it were in normal conditions it would reliably promote true beliefs when operating according to the design plan. A brain-in-a-vat might operate according a good design plan, and so enjoy the first level of warrant. It's at this level that his use agrees with mine. The second requires being in normal conditions too, and so requires globally reliably producing true beliefs. The third requires the absence of 'local relevant alternatives' like fake

barns, and so requires 'local' reliability. When all three are in place, 'warrant' entails truth. And so when his followers say that warrant entails truth, it's the positive epistemic status that involves all three levels that they are talking about. I compare and contrast his view with mine elsewhere (2011c; in preparation-a). Since his first level accords with mine, he too would recognize the role, despite his focus on knowledge.

An objection Plantinga and many others would lodge against my account is that it wrongly assumes that Mother Nature cares about truth. She only cares about what contributes to survival and doesn't care one whit about truth. So long as it works, it doesn't matter if it gets things right. Though there's an obvious truth in this point, it does not touch my view. The obvious truth is that it's not *a priori* that natural selection would select for true beliefs or reliable belief-forming systems; it's not *a priori* necessary that any representational system have the evolved or learned function of getting things reliably correct. But a representational system could still, empirically, have, as a matter of fact, been selected for reliably getting things right. Just as it's not *a priori* that Mother Nature select for hearts, but it's empirically entrenched that she selected for them nonetheless; it's empirically well-established that Mother Nature selected our eyes, among other things, for reliably informing us about the properties and attributes of distal objects in our natural environments. I discuss these issues at greater length elsewhere (Graham 2011a; 2011b).

I have sketched distinctions between different levels of mind or psychological capacity out of the conviction that different epistemic statuses sometimes correlate with different psychological capacities. This is surely true of being justified as I defined it; it's a status that requires the capacity to exercise critical reflection. It's also true of warrant. For warrant arises with perception and the capacity to reason, but doesn't require the capacity to critically reason. That's not to say that beliefs formed and sustained through critical reason lack warrant, but only to say that warrant emerges prior to the development and exercise of critical reasoning.

My methodology focused on carving out a role for a distinctive epistemic status. I've relied on a four-level account of psychological kinds, the general idea of a positive epistemic status, and a general theory of functions and goods in order to isolate both a distinctive epistemic role and a realizing property. You can accept the role without accepting the realizing property.

1 I didn't study much epistemology in graduate school, though I wrote a dissertation on testimony and soon found myself teaching epistemology courses and seminars. Robert Audi's collection of papers in epistemology was a godsend. And more than any of the others, 'Justification, Truth and Reliability' made a huge impact on my understanding of the central issues in epistemology, especially the theory of justification. I had the further good fortune early in my career to get to know Robert personally. I continue to benefit from his advice, encouragement, and philosophical good sense.

2 Sellars is rightly accused of denying representational capacities to animals in "Empiricism and the Philosophy of Mind." Sellars (1981) later cautioned against denying animals genuinely perceptual representational systems. Animals, he claimed,

possessed 'propositional' representations but lacked 'logical' representations. Propositional representations refer and predicate, while logical representations require the capacity to draw the whole suite of logical inferences, and so require concepts for logical concepts, variables and quantifiers. As I am using the terms, a representation with non-propositional structure refers and predicates but does not require the whole logical suite, but a representation with propositional structure requires at least some logical inferential capacities, though not necessarily all. An animal might be able to perform disjunctive syllogism, for example, without being able to represent generality. Logical capacities are possibly diverse, and needn't come as a complete package. So a Sellarsian 'propositional' representation can have a non-propositional structure, and representations with propositional structure requires logical concepts and capacities, but not necessarily the whole suite; propositional representations needn't be fully 'logical' representations.

3 Warranted *reasons* are propositional contributions to warrant. When we reason from one belief or set of beliefs to another (from premises to conclusion), our propositionally structured beliefs comprise our *reasons*. Reasons are, or play the role of, arguments. The premises, when warranted, contribute to the warrant for the conclusion.

Reasoning transitions from premises to conclusion. When we reason well we conform in our reasoning to various rules or forms of good inference: so conforming contributes to the warrant for the conclusion; not just any old transition will do.

We do not always represent the rules of good inference or that fact that our reasoning is good in a good amount of our reasoning—especially unreflective, uncritical reasoning. We are not always aware, or do not always appreciate, that we are making a good inference. We don't have as a premise within the argument the recognition that our reasons and reasoning conforms to a good inference rule. Nor do we have a recognition or awareness outside of the argument that our reasons and reasoning conforms to a good inference rule. Critical reasoning differs. In critical reasoning—especially reflective reasoning—we are usually aware that our reasons and reasoning are good. 'Inferential internalism' may be true of critical reasoning (Leite 2008).

So in ordinary, unreflective, uncritical reasoning we are also entitled to rely on transitions from premises to conclusion. We are entitled to rely on forms or patterns of good inference. The premises and the entitlement to rely on the transition from premises to conclusion contribute to the overall warrant for the conclusion. The warrant for beliefs through inference is a mix of both warranted reasons and entitlements, while the warrant for perceptual beliefs is paradigmatically just from entitlements.

I failed to notice some of these distinctions when I wrote "Testimonial Entitlement and the Function of Comprehension" (Graham 2010b). In that paper I claimed that comprehension-based beliefs enjoy entitlements for a function of comprehension-with-filtering is to reliably induce true beliefs. I failed to note that comprehension-states—our representations of the assertive speech acts of others—have propositional structure. I now see that our warrant for comprehension-based beliefs is mixed: it involves reasons and entitlements. What matters to a broadly anti-reductionist point of view regarding testimony is that not all of our warrants for testimony reside solely in reductive first-hand arguments for the reliability of our interlocutors.