

CHAPTER 2

Warrant, functions, history

Peter J. Graham

1 Virtue epistemology and proper function

According to John Greco and John Turri, virtue epistemology has two basic commitments.¹ The first is that epistemology is a normative discipline and not merely a branch of natural or social science, *pace* Quine. This implies, among other things, that epistemology should focus on epistemic norms, epistemic value, and epistemic evaluation.

The second is that epistemology should follow a direction of analysis found in virtue ethics. Broadly speaking, in virtue ethics the moral rightness of an action is determined by the properties of the agent; the action is right only if based on the ethical virtues of the agent. Similarly, in virtue epistemology the epistemic rightness of a belief is determined by properties of the agent that caused the belief; the belief is warranted or knowledge (or otherwise epistemically valuable) only if based on the epistemic virtues of the agent.

Virtue epistemology standardly divides into two camps: virtue-reliabilism and virtue-responsibilism. The two camps talk about different things when they talk about epistemic virtues. Virtue-reliabilists talk about reliable belief-forming faculties such as perception, memory, and reasoning. Virtue-responsibilists talk about character traits such as open-mindedness and conscientiousness.

I am a virtue-reliabilist about epistemic warrant, for I seek to understand epistemic warrant in terms of features of reliable belief-forming faculties.

I presented an earlier version of this chapter to an audience at Soochow University in Taiwan. I am grateful to feedback on that occasion, and especially from Ernest Sosa. I am grateful for useful comments from Colleen Macnamara and Zach Bachman. This work was supported by a research grant from the UC Riverside Academic Senate and by a Visiting Professor Fellowship at the Northern Institute of Philosophy, University of Aberdeen.

¹ Greco and Turri 2011.

Some philosophers use ‘warrant’ for that property that converts true belief into knowledge. I do not use ‘warrant’ this strongly. Instead I use ‘warrant’ the way most epistemologists use ‘justification’ or ‘justifiedness’. Warrant is then not sufficient for converting true belief into knowledge for warrant so understood does not metaphysically entail truth. I prefer ‘warrant’ because ‘justification’ connotes the ability to justify, and so tends to over-intellectualize knowledge, especially perceptual knowledge and the knowledge of small children and higher non-human animals.²

I believe the best way to develop virtue-reliabilism about warrant is along proper functionalist lines, for virtue is a teleological notion. An epistemically virtuous process is a normally functioning belief-forming process that has forming true beliefs reliably as a function; virtue-reliabilists should be proper functionalists.

I have argued elsewhere that epistemic warrant consists in the normal functioning of the belief-forming process when the process has forming true beliefs reliably as a function.³ A belief is warranted when the process has functioned normally, when normal functioning is constitutively associated with reliably inducing true beliefs and avoiding error.

The adequacy of a proper function virtue-reliabilism, however, obviously turns on the nature of functions. I endorse the *etioloical* theory of functions associated with Larry Wright, Ruth Millikan, Peter Godfrey-Smith, Karen Neander, and many others. On the theory, functions turn on histories that explain why the item exists or operates the way it does. If warrant requires functions, and functions require history, then warrant requires history.

At least within epistemology, it is commonly thought that if warrant requires functions, then warrant requires natural selection. And so many within epistemology are inclined to see my view as requiring a history of natural selection over generations for warrant.⁴ Indeed, according to the philosophical zeitgeist, if you hold a “proper function” view of warrant, you think only God or Mother Nature can assign functions.⁵ So if

² For more on my use of ‘warrant’, see Graham 2011a. See also Burge 2003.

³ Graham 2010, 2011a, 2011b, 2011c, 2012a.

⁴ Sandy Goldberg writes “Most recently Peter Graham [has] endorsed the view that only ‘naturally evolved’ processes count [towards doxastic justification] ... [where a process confers justification iff reliable] in those environments in which the process evolved” (2012: 109, 117). David Copp says that proper functionalism “rests on the claim that the human cognitive system was selected for in the evolutionary development of the species” (this volume). See also the friendly presentation of my view in Evans and Smith 2012: 194 and the implicitly critical one in Lyons 2011. I am partly responsible for this reading of my work, having emphasized natural selection as a source of function for perception and our capacity to comprehend and filter assertive communication.

⁵ Witness Richard Feldman and Andrew Cullison: “According to the proper functionalist view about justification, justification is a matter of forming beliefs in accordance with a design plan that is

you think God did not design us, then you must think Mother Nature did all the work. So if warrant requires functions, then warrant requires generations.

But isn't that asking too much? What about learning? Can't learned perceptions and acquired belief-forming competencies warrant their corresponding beliefs? Why *nature* and not also *nurture*? And what about Swampman? Can't he have warranted beliefs, even if he has no history at all?⁶

I shall argue that proper functionalism without God does not require natural selection per se for functions, and so does not require natural selection per se for warrant, for natural selection is not the only source of etiological functions. I discuss sources of functions that take considerably less time than natural selection. If warrant turns on functions, and functions turn on history, then warrant turns on history. How much? In some cases, not much.

By treating epistemic virtue in terms of functions, and functions in terms of history, have I set out to *naturalize* epistemic virtues, and so to naturalize epistemic warrant? No. Rather I have set out to *understand* functions, virtues, and warrant. If it turns out, however, that the account of functions I advance is naturalist in the intended sense, then my account will satisfy those with a naturalist agenda.⁷

aimed at acquiring true beliefs. While the view mentions a *design plan*, the view is supposed to be theologically neutral. That design plan may come from God (if there is one) or via natural selection and evolution (if there is no God)" (2012: 98). In Sosa's discussion of Plantinga, he too suggests that a design plan (and so a function or purpose) can arise only from God or Mother Nature. "The problem for proper functionalism is [that it takes] it to be impossible that there be someone with warranted belief who has no design plan imposed by any agency or process that designed him ... [for proper functionalism requires design by God or evolution for warrant]" (1993: 55, 57).

⁶ Goldberg, Copp, and Lyons reject proper functionalism for this reason. Goldberg and Lyons think Swampman would have justified beliefs, without saying why Swampman would have beliefs, let alone warranted beliefs. Copp says epistemology should ignore our evolutionary roots. "Perhaps human beings popped suddenly into existence, out of thin air ... Nothing in epistemology turns on whether our cognitive system is the way it is due to its having been selected for" (this volume). Sosa (1993) imagines a Swampbaby that is discovered by a hunter and then grows to adulthood, having learned all sorts of things. He proposes this as a counter-example to proper functionalism.

⁷ According to Neta 2007, the account I offer here counts as naturalist, or nearly naturalist. Kornblith 2002 would probably agree. For a structurally similar view of warrant that ties warrant to functions and functions to norms, see Burge 2003. Burge ties warrant to representational functions. Burge believes it is a priori necessary that the function of the perceptual system is to perceive, and to perceive is to represent accurately. Hence the representational function of the perceptual system is to represent accurately. He argues similarly for any belief-forming competence or capacity. He also argues that representational functions do not reduce to biological (or other obviously "naturalistic") functions. I discuss Burge in Graham *in press a* and *in press b*.

The rest of the chapter is organized as follows. In [sections 2](#) and [3](#) I explicate etiological functions. You'll see why etiological functions require history. In [sections 4](#) and [5](#) I review my proper functionalist view of warrant. Then in [sections 6](#) to [10](#) I discuss sources of functions.

2 Etiological functions

Like many words, 'function' has many overlapping and related senses. In the sense I intend, the *function* of a thing denotes what it's for, its purpose. The heart is supposed to pump blood; that's its function; that's what it's for.⁸ Functions in this sense are (typically) effects. By beating the heart causes the circulation of blood. But not every effect (even highly regular effects) is a function in this sense. Your heart regularly and reliably makes a rhythmic noise, but making noise is not a function of your heart; that is not what it is for. Your nose regularly and reliably holds up glasses, but you do not have a nose in order to hold up glasses. There are functional effects that explain why something exists, and then there are non-functional, "accidental" side-effects that do not.

Larry Wright (1973) argued that this distinction strongly supports an *etiological* condition on functions, where functions are consequences that *explain* why the item exists. Here is Wright's analysis:

A function of X is Z if and only if:

- (1) X does Z (Z is a consequence [result] of Xs being there, i.e. Xs are disposed, do, or can do Z).
- (2) X is there because it does Z (that Xs are disposed, do, or can do Z explains why X is there).

Wright's condition (2) then says that for any function, there must be some *feedback mechanism* that takes the satisfaction of (1) as input and generates existence or continued existence as output. Functions thus arise from *consequence etiologies*, etiologies that *explain* why something exists or continues to exist in terms of its consequences, because of a feedback mechanism that takes consequences as input and causes or sustains the item as output. Functions are then *explanatory* features or effects.

Non-functional features or effects are *non-explanatory* features or effects, and so in that sense "accidental," even if non-accidentally regular. By

⁸ There are some so-called philosophical naturalists who deny that the heart has a function in this sense, for they deny *functions* in this sense, and certainly deny functional *norms* in this sense. I shall discuss such views in future work.

beating regularly, hearts pump blood, and we have hearts because they pump blood. Though by beating regularly hearts make noise, we do not have hearts because they make noise. Noses keep air warm and dry, and we have noses because they keep air warm and dry. And though they hold up glasses or nose rings, we do not have noses because they hold up glasses or nose rings.

Malformation raises an obvious difficulty. Consider a heart that's heavily malformed. Malformed, it can't pump blood, and *it* certainly does not exist because *it* can pump blood. But then *this heart* fails both of Wright's conditions. Even so it still has the function of pumping blood; that's what it is supposed to do.

This difficulty is easily avoided by incorporating a type–token distinction. True, certain malformed token hearts cannot pump blood. But the *type* can have the function of pumping blood provided a feedback mechanism takes past *token* hearts as input and produces or maintains hearts *because* past token hearts pumped blood. Then the heart (type) exists because it (tokens of the type) pumps blood.

Distinguishing types from tokens has other benefits. For instance, you may have wondered how future consequences can explain present existence. We can dispel the worry, for *current* tokens of functional items acquire their functions from *past* tokens going through a feedback mechanism. The past thus explains the present.

There's another problem not so easily avoided. Mark Bedau (1991) noticed that Wright's definition applies to some cases of non-living, inorganic materials. He describes a case from Richard Dawkins involving clay crystals that build dams in streams. The dams result from layers of sediment stacking up on top of one another according to the pattern laid down by the crystal. As a result the crystals replicate themselves; the dam is a tower of new crystals. Once the dams are built, the stream cannot wash the clay downstream; the dam thus stays in place. These crystals build dams, and they exist because they build dams; the crystals meet Wright's two conditions. But intuitively building dams is just something these crystals *do*, not something they are *supposed* to do. Intuitively there's *nothing* they are for; they have no purpose or function.

Following Hempel (1959), Ruse (1971), and Bedau (1991), Peter McLaughlin (2001) argues that we should include a *benefit* or *welfare* condition. Functions are not just explanatory features or effects. Functions are means to some *good* or *benefit* of the containing system. In order for Z to be a function of X, doing Z must do the system of which it is a part some good, and this good must be relevant to the feedback mechanism

that explains why X exists in the system. Functions arise through a feedback mechanism that involves explanatorily *beneficial* effects. Pumping blood helps *you* survive; pumping blood is a means to many of your ends; it clearly does *you* a lot of good. That's why, according to Hempel, Ruse, Bedau, and McLaughlin, it's a function of your heart to pump blood. The clay dam, on the other hand, doesn't have a good. Replicating isn't a means to any end, for the crystals or the dams have no ends, either as individuals or as members of a kind.

Putting Wright, Hempel, Ruse, Bedau, and McLaughlin (2001: 83) together, we arrive at the following abbreviated analysis of natural functions:

A function of X in S is Z iff:

- (1) X does Z in S.
- (2) Z benefits S.
- (3) X exists in S because Z benefits S (X is the product of a feedback mechanism involving the beneficial character of Z to S).

In this account, condition (3) says that for any function, there must be some feedback mechanism that takes the satisfaction of (1) and (2) as input and generates existence or continued existence as output. In recent work Larry Wright agrees: consequence etiologies that ground functions are *virtue* etiologies (Wright 2012).⁹

There are two points deserving emphasis. The first is that the whole purpose of turning to history is to mark the difference between explanatorily beneficial effects and other, non-explanatory, accidental effects, whether beneficial or not. The turn to history is driven by the need to distinguish explanatory functional from non-explanatory accidental effects. Without any history, there's no basis to draw the distinction.

The second is that the etiological account of function, as stated, is entirely neutral on possible feedback mechanisms. The account does not entail any particular feedback mechanism; it allows any possible feedback mechanism to generate functions, as long as it is a feedback mechanism taking beneficial characters or effects as input and produces existence or persistence as output. Etiological functions metaphysically entail feedback; they do not metaphysically entail any particular kind. This is what allows for a plurality of actual sources of etiological functions.

⁹ Adding the beneficial effects condition marks a change in my view, for I did not include it in earlier work. It also makes naturalizing functions, and so naturalizing warrant, more challenging, for now the notion of benefit must be explicated. I will not try to do so here.

3 Normal functioning

The etiological account of functions entails an account of normal functioning and normal conditions. On the etiological account, functions arise when an item produces a beneficial effect that in turn enters into a feedback mechanism, where the mechanism explains why the item persists or reoccurs because of the beneficial effect. The full explanation for why and how all of this happened will cite how the item *worked* or *operated* so as to produce that effect and the *circumstances* – both internal or “inside” and external or “outside” the individual or organism.

What counts as normal functioning and normal conditions falls out of the historical explanation. *Normal functioning* is the way the item *worked* or *operated* when it underwent feedback for its beneficial effect; *normal working just is* working that way. Normal conditions are those circumstances (and circumstances of relevantly similar kind) where all of this happened. Look at the item’s history, at the beneficial effects that help explain why it persists and recurs, at how it worked to produce these effects, and where it all happened. *Voilà*, normal functioning and normal conditions (Millikan 1984).

For example, a muscle in an organism’s chest pumps blood by beating regularly. In turn it is connected in a systematic way with other parts of the organism, embedded in a certain type or kind of environment. If pumping blood explains, in part, why the muscle recurs through benefiting the kind or the individual, then it comes to have pumping blood as a function. The way the muscle worked when it entered the feedback mechanism for pumping blood equals normal functioning. Normal conditions are then those circumstances (and circumstances of similar type) where all of this occurred.

Given the way normal functioning and normal conditions are determined, normal functioning and normal conditions are then constitutively, explanatorily interrelated with function fulfillment. Normal functioning, normal conditions, and function fulfillment are all holistically interrelated. In particular, *normal functioning* is individuated and explanatorily understood in terms of the function of the item, for *normal functioning just is operating or working* the way the item operated in normal conditions so as to produce the functional effect. *Normal functioning* is then constitutively associated with function fulfillment.

Normal functioning constitutively “aims” at, contributes to, and conduces function fulfillment. For normal functioning is non-accidentally and explanatorily understood in terms of the function (and so the “aim”)

of the item. By functioning normally, the item non-accidentally and constitutively fulfills its function (and so achieves its “aim”). By functioning normally, it non-accidentally and constitutively contributes to function fulfillment; normal circumstances contribute the rest. And by functioning normally in normal conditions, it non-accidentally and constitutively conduces function fulfillment.¹⁰

Though holistically interrelated, it’s important to see that normal functioning and function fulfillment are token-distinguishable; on particular occasions you can have one without the other. Consider a world-famous surgeon who needs to remove your heart during a very complicated surgery to cure a disease in the middle of your chest. She may place your heart in a sterile dish and stimulate it with electrical wires so that it beats normally – it operates exactly the way it should – but no blood is passing through. Your heart then functions normally (it’s in perfect shape), though it doesn’t fulfill its function. And so on occasions a normally functioning heart may fail to fulfill its function for it’s not in normal conditions.

4 Three functional norms

Before turning to known feedback mechanisms on beneficial effects, I will review my view of warrant as turning on functions. After all, without seeing why warrant should turn on etiological functions, one might wonder what the fuss is all about. But first I shall spell out the category of functional norms, for I see the normativity constitutive of warrant as a species of functional normativity.

When philosophers discuss norms, they typically discuss norms that prescribe or guide behavior or thinking, such as prudential, moral, or social norms. Prescriptive or guiding norms are norms that we can represent, discuss, consider, internalize, subscribe to, consciously follow, flout, debate, challenge, and so on. Social norms, for example, are regularities in behavior in a group prescribed by members of the group. If the norm were not represented by a sufficient number of the group, the group would not prescribe the behavior, and so the norm would not exist.

There are norms that do not require the capacity to represent, think, internalize, or subscribe to a norm. There are norms that are neither prescriptions nor guides. Functional norms are norms in this broader sense.

¹⁰ This marks an interesting difference with functions from conscious, intentional design. An item can have a function from conscious, intentional design without ever fulfilling its function. Just think of the dustbin of failed inventions. Thus whatever fixes normal functioning and normal conditions isn’t holistically interrelated with function fulfillment for consciously assigned functions.

When there are functions, there are norms. Functional norms are standards or levels of “possible performance that is in some way adequate for fulfillment of a function or purpose.”¹¹ The heart’s function, then, determines standards or levels of performance in fulfilling its function. Given the heart’s function, we can ask how well it performs. When it fulfills its function and operates normally, it meets levels of adequacy for performance in fulfilling its function and thereby fulfills norms. Such norms need not prescribe or guide. No norm tells the heart what to do. The heart does not look up or represent any norm to guide its activity. Functional norms are a broader kind than prescriptive or guiding norms. They do not depend on the aims or intentions of individuals, on being represented or being endorsed.

I identify three functional norms for any item with an etiological function: function fulfillment, normal functioning, and function fulfillment *because* functioning normally. Function fulfillment is trivially a level of performance adequate for fulfillment of a function. When your heart pumps blood, it meets a norm trivially associated with its function. Normal functioning is also a level of performance in some way adequate for fulfillment of the item’s function, for it is the explanatorily relevant way the item non-accidentally fulfills its function in normal conditions; it is the way the item is *supposed* to *work* or *operate* so as to fulfill its function. When your heart functions normally, as it should, then it meets a second norm associated with its function. Function fulfillment *because* functioning normally is likewise meeting a norm in this sense; it is the explanatorily relevant way the item fulfills its function in normal conditions through meeting norms adequate for the fulfillment of its function or purpose.

Notice the structural parallel with Ernest Sosa’s three-part normative structure for the exercise of a competence, a structure also elaborated by Wayne Riggs, Duncan Pritchard, and John Greco, among others. For Sosa, a competence is a reliable capacity to achieve some *aim*. Aim stands to function. If the competence fulfills or achieves its aim, the exercise of the competence is *accurate*. Accuracy stands to function fulfillment. If the competence is reliable in normal conditions, and the exercise is non-defective, then the exercise of the competence is *adroit*. Adroitness stands to normal functioning. If the exercise is also accurate because adroit, then the exercise is *apt*. Aptness stands to function fulfillment because functioning normally.

¹¹ Burge 2010: 311. I have taken the contrast between functional norms and prescriptive and guiding norms from Burge. Burge calls functional norms *natural* norms. Though natural, he does not intend to convey that they are naturalistic in some strong, reductionist sense of the term.

If aims require the individual who possesses the competence to represent the aim, then the category of functional norms and functional achievements is broader than Sosa's three-part structure, for functional norms do not require, as such, any representational capacities. Functional norms arise prior to mind. But since aims can determine functions, given a feedback mechanism, Sosa's three-part structure could be a species of functional normativity. Functional normativity does not exclude represented aims, goals, and so on.

For norms for items with etiological functions, normal functioning *encodes*, for partly constituted by, function fulfillment. We then understand the second norm of normal functioning in terms of the first norm of function fulfillment – the second norm is constitutively associated with the first – for normal functioning is constitutively associated and explanatory understood in terms of function fulfillment for items with etiological functions.

5 Epistemic warrant as normal functioning

I now turn to my account of warrant. Assume a belief-forming process has forming true beliefs reliably as an etiological function. There are then three functional norms it can meet: function fulfillment, normal functioning, and function fulfillment because normal functioning, where the second (and so the third) is constitutively associated with the first. Since these norms are understood in terms of promoting true belief and avoiding error, they are epistemic norms. Some epistemic norms are then functional norms. Meeting these norms are then epistemic achievements, goods, or successes.¹²

Epistemic warrant consists in fulfilling epistemic norms, for warrant is a normative status or achievement. Epistemic norms are norms understood in terms of promoting true belief and avoiding error. But this premise does not itself establish that warrant consists in functional normativity, for there are many kinds normativity, and so possibly many kinds of epistemic normativity.

There are, for instance, epistemic norms that prescribe and guide. Some tell us when we should inquire and for how long. Some prescribe techniques of critical reflection; they tell us how to assess and evaluate reasons

¹² This argument is independent of my view of warrant. You can deny that warrant requires functions and still accept that functions entail norms, that there are epistemic functions and so epistemic norms associated with those functions, and so accept that there are epistemic achievements, goods, or successes that consist in fulfilling functional norms.

for and against. Some prescribe thresholds for decision, conviction, or judgment. Some epistemic norms also guide. They guide when we inquire and for how long. They guide our reflection. They guide our judgments and levels of conviction. When they guide, we represent, endorse, and follow these prescriptions. We consult the norms in order to guide our inquiry, reflection, and assent.¹³

But I do not believe warrant requires following norms that prescribe or guide. This is partly because I believe warrant applies broadly throughout the animal kingdom. Higher non-human animals, small children, and even ordinary adults either cannot, or need not, represent or think epistemic prescriptions or epistemic guides for their beliefs to enjoy warrant. A chimp or ape can have a warranted perceptual belief without even having the capacity to think about functions, norms, truth, accuracy, warrant, evidence, thresholds, and so on. Since warrant consists in fulfilling epistemic norms, but warrant applies broadly, the relevant norms cannot be prescribing or guiding norms. The normativity constitutive of warrant is neither prudential, moral, nor social.

I believe warrant consists in fulfilling functional norms associated with the epistemic function of the belief-forming process. Of the three functional norms, I believe warrant consists in fulfilling the second, the norm of normal functioning.

I believe this because I believe warrant may persist outside of normal conditions, and a system may still function normally outside of normal conditions, though it normally cannot explanatorily and non-accidentally fulfill its function outside of normal conditions. If warrant required function fulfillment, then warrant would be restricted to normal conditions.

An individual may stumble outside of normal conditions without any awareness that this is so. The animal's functional capacities may still function normally, even though the individual is no longer in normal conditions. The animal's capacities then fulfill a norm constitutively associated with function fulfillment. This is then a functional good or success. Organs too may operate normally outside of normal conditions. Recall the heart removed from the chest during a complicated surgery; it may operate (beat) normally, even though it is not pumping any blood, for it is no longer in normal conditions.

¹³ I have discussed epistemic norms that prescribe and guide in Graham [in press c](#). In a similar vein, Copp (this volume) isolates epistemic norms that are social norms that approximate ideal norms, where ideal norms are those that, when followed, solve our needs to acquire relevant information to help meet other needs. Kornblith (2002) identifies epistemic normativity with instrumental normativity. Though these are all genuine kinds of epistemic normativity, I do believe warrant involves a different kind of epistemic normativity.

If the function of a belief-forming process is an epistemic function, then functioning normally is an epistemic good or success, even if the individual or the process is no longer in normal conditions. So there are epistemic goods that persist outside of normal conditions. The massively deceived, disembodied brain-in-a-vat vividly illustrates such a possibility. Outside of normal conditions, without any awareness that this is so, an individual's belief-forming processes may function normally. The massively deceived brain-in-a-vat may still function normally, despite failing to fulfill its epistemic functions.

This view of epistemic warrant as normal functioning is reliabilist in spirit, for it sees warrant as constitutively associated with promoting true belief and avoiding error, for normal functioning is constitutively associated with reliably getting things right when that is the etiological function of the belief-forming process. Warrant entails reliability in normal conditions, for warrant is grounded in reliability in normal conditions. Warrant, however, is not restricted to normal conditions.

This view explains why warrant aims at, contributes to, and conduces truth, three traditional marks of epistemic warrant. Normal functioning "aims" at reliably getting things right, and so "aims" at truth. Normal functioning non-accidentally and explanatorily contributes to reliably getting things right. And in normal conditions, normal functioning non-accidentally and explanatorily conduces towards true belief.

If warrant is constitutively, explanatorily, and non-accidentally associated with promoting true belief and avoiding error, if warrant consists in meeting epistemic norms, if animals and small children can meet those norms, and if warrant persists outside of normal conditions, then warrant consists in normal functioning when the belief-forming process has forming true beliefs reliably as an etiological function.

Since functions require history, warrant requires history. But if warrant requires history, how much history does warrant require? It all depends on how much history etiological functions require. In the remainder of the chapter I review sources of functions that don't require an awful lot. I'll start, however, with natural selection, which seems to require an awful lot.

6 Directional and maintenance selection

Natural selection requires three elements: variation, copying (inheritance), and beneficial consequences (fitness). Imagine birds that use their color vision to prey on a population of beetles. Imagine these beetles vary in

color: half are brown and half are green. Imagine further that the beetles feed and live on a leafy green plant. The green ones are hard to see and so more likely to live long enough to reproduce. The brown ones, on the other hand, are easy prey. Now assume that their coloration is inherited, so that green beetles are more likely to produce green offspring than brown, and brown are more likely to produce brown offspring. Over time green coloration will come to predominate in the population. And now we have change in the population of beetles over time: once fifty–fifty, nearly all are now green. Within the beetle population coloration *varies*; their coloration is *inherited*; and coloration has obvious *consequences*: green beetles are camouflaged in their natural habitat; brown beetles don't stand a chance. Here's a case of modification with descent – evolution – through natural selection.

Natural selection takes time; it works over generations. As a feedback mechanism it takes frequencies of beneficial traits in earlier generations as input and produces frequencies in later generations as output. Even so, evolution by natural selection can happen very fast for organisms that reproduce rapidly (think of fruit flies and bacteria). But for organisms like us, evolution by natural selection often moves very slowly.

Or at least this is obviously true for *directional* selection, selection that leads to a change in the frequency of certain traits. *Maintenance* selection, on the other hand, maintains the frequency of traits in a population. Most mutations, for example, are harmful. They produce malformations that often lead to death well before the opportunity to reproduce, or diminish opportunities for reproduction. Because harmful they are selected out, in favor of the normal variant of the trait. The non-malformed trait then continues to exist and predominate in the population because of its relative superiority. Because of the prevalence of such harmful mutations, nearly every trait in a population is currently undergoing some form of maintenance selection.

Maintenance selection is full-blooded natural selection. It involves variation, inheritance, and beneficial consequences. It explains why a type of trait is preserved, upheld, or maintained in a population. And so it assigns functions for the very same reason that directional natural selection does. Maintenance selection, like directional selection, is a feedback mechanism on beneficial consequences.

It is a contingent, empirical question how many generations are required for directional selection to assign functions, or a change of functions. If the human heart last underwent directional selection 10,000, 30,000, or 100,000 years ago for its current form and function, then

functions from directional natural selection would require considerable history indeed. But since the human heart undergoes maintenance selection in every generation, functions from maintenance selection require considerably less. If humans were, *per impossible*, created out of nothing a generation ago, maintenance selection would be at work, assigning functions.

7 Self-replication and repair

I now discuss another source of etiological functions. In *What Functions Explain*, McLaughlin argues for a non-hereditary feedback mechanism alongside natural selection (2001: 162–90). He thinks the ordinary metabolic activity of an organism that sustains the organism's own self-replication and repair fits the bill.

The ordinary operation of your metabolism keeps you alive. And this is partly because the ordinary metabolic activity of your systems and subsystems involves repairing and replacing the cells of your body, and so continually repairing and replacing the various systems and subsystems that make up your body. Your heart, by pumping blood, contributes to its own reassembly and repair, and thereby keeps you alive. Your heart thus persists in your body because of a feedback mechanism – normal metabolic activity – that takes earlier cycles comprising one group of cells as input and produces later cycles comprising another group of cells as output. An earlier *cycle* of your heart, by pumping blood, contributes to the existence of a later *cycle*, partly in virtue of its beneficial effects to you. Any trait that is advantageously integrated into the normal metabolism of your body contributes to *its* own reassembly and repair, and thereby to your continued life.

How does your metabolism generate functions? Once again, take your heart. The normal operation of your metabolism generates pumping blood as function provided (1) your heart pumps blood; (2) pumping blood (and so blood circulating through your body) benefits you; and (3) your heart exists or persists in your body because pumping blood benefits you. Both natural selection and your metabolism generate the same function for your heart. It's then over-determined that a function of your heart is to pump blood. A belief-forming psychological capacity will then have the function of reliably inducing true beliefs via the creature's metabolism provided (1) it reliably induces true beliefs; (2) reliably inducing true beliefs benefits the creature; and (3) the creature possesses such a capacity because reliably inducing true beliefs benefits the creature.

The normal operation of the metabolism of an organism is a feedback mechanism alongside natural selection taking earlier beneficial effects as input that produces functions as output. Where natural selection takes earlier generations (and so distinct tokens of the type) as input, the normal operation of the metabolism of an organism takes earlier cycles as input (and so maintains a token of the trait through time); it explains the continued existence of traits within individuals over *cycles* in virtue of their beneficial effects.¹⁴ A first-generation trait within an individual organism could then acquire a function provided it is advantageously integrated into the metabolism of the organism.

This difference suggests a difference between the *species* or *population* as system and the *individual* as system. By pumping blood, hearts contribute to the survival of individual organisms long enough to reproduce and propagate the species or population. That's clearly a benefit to the species or the population. And by pumping blood, hearts contribute to the survival and wellbeing of the individual organisms themselves, which is clearly a benefit to the very individual in question, never mind the species or population.

These two mechanisms interact. Natural selection benefits the individual by preserving traits beneficial to individuals; you have your beneficial traits because of a long ancestral history involving natural selection. And normal self-replication and repair benefits the species or population, for if you don't self-replicate and stay alive, you can't propagate your kind. But they do not always overlap, which shows why we should distinguish the two. For many creatures engage in activity that benefits only the species. Salmon swim upstream to fertilize eggs only to die. Some male spiders, right after mating, get immediately killed and eaten by the female. Some creatures hatch their eggs internally. The hatchlings then eat their way out, obviously killing their mother in the process. Mules provide an example of the opposite kind of case, where their organs clearly benefit the individual mule, but nothing they do contributes to reproduction, for mules are sterile.

8 Interlude on Swampman

This distinction between *cycles* and *generations* – between metabolism and natural selection as mechanisms – helps dispel Swampman, for

¹⁴ And so we need to qualify our earlier point about the importance of the type–token distinction in our account of functions: functions require either earlier *tokens* of the item for the type to acquire a function, or earlier *cycles* of the token for the token to acquire a function. But even then we'll still need the type–token distinction to accommodate malfunctioning, for a malfunctioning token may never fulfill its function; earlier cycles then don't contribute to later cycles of that very token.

metabolism takes very little history to generate functions. Swampman is a creature of philosophical science fiction. Imagine a bolt of lightning hitting a log in a swamp and creating a molecule-for-molecule duplicate of Barack Obama, a duplicate that bears absolutely no causal or explanatory connection to the real Barack Obama or to any other real human being, living or deceased. The physical duplicate – Swampman – is then a cosmic accident of vast proportions. Assuming for the sake of argument such a possibility, some philosophers find it natural to say that Swampman’s “heart” has a function just like Obama’s heart, even though Swampman bears no causal or explanatory relation to Obama, or to any other real human being or biological entity, living or deceased.

Swampman is a full-body example of what biologists call hopeful monsters. In nature, new organs or traits often arise very slowly through a series of micro-mutations; small changes through a gradual process of variation, selection, and replication. But once in a blue moon a macro-mutation arises: an almost entirely new trait or organ, very different in kind from its ancestral trait. In actual cases nearly all of these are deleterious to the recipient; the recipient soon dies or is unable to reproduce and the trait is selected out through maintenance selection. Think of extreme birth defects. But sometimes one of these macro-mutations actually benefits the recipient. These traits are called “hopeful monsters.” A hopeful monster is a beneficial macro-mutation. Since they are mutations, they don’t have an evolutionary history; they are “first-generation” traits. Thus they don’t exist because of natural selection; natural selection works only on traits that already occur (though it can increase the probability that various traits will emerge). Swampman is just this sort of case taken to the extreme.

If Swampman’s organs have functions, must we reject the etiological account? Hopeful monsters are a problem only on the assumption that natural selection is the only feedback mechanism generating functions. But since it’s not, hopeful monsters are not a problem. True, the first cycle of a hopeful monster has no function, even if it has a beneficial effect. But without a feedback mechanism in play there’s no distinguishing between functional effects and merely accidental, albeit beneficial, effects; only persistent and recurrent traits have functions (Hempel 1959; McLaughlin 2001: 67–68, 168); that’s the whole point of the function–accident distinction. Once the hopeful monster starts to benefit the organism and thereby contribute to its own self-replication and repair, it enters a feedback loop that partly explains its own continued existence. Its effects are then functional, not merely accidental, for its effects play an explanatory role. The

same holds for all of Swampman's beneficial organs and traits. Swampman at the moment of his creation has no functions; over time Swampman's organs and traits acquire functions. But then we have an explanation for why someone might think his organs do indeed have functions, for over time they do. The existence of hopeful monsters – even Swampman – is not a problem for an etiological theory of functions that takes a broader view on feedback mechanisms. Some history is required for functions, just not an awful lot.

Of course the recalcitrant philosopher may insist that Swampman has functions at creation. If they do, all I can do at this point is pass the baton, and invite them to develop a better account of functions, an account that treats the function–accident distinction just as well without any appeal to any history whatsoever, consistent with their recalcitrance. Good luck.¹⁵

9 Trial-and-error learning

I now turn to a third feedback mechanism: trial-and-error learning. Psychologists call trial-and-error learning *operant conditioning* or *instrumental learning*. Imagine a four-year-old learning to tie his shoes. As any parent knows, this isn't a trivial task. My parents moved my fingers for me. I got one part of the process partially right, but the rest was a mess. Somehow, over time and with enough effort, I learned to tie my shoes. Now it's effortless.

It's a trite observation in textbooks on learning and memory that trial-and-error learning parallels natural selection. Natural selection requires variation, consequences, and copying: the variation is genetic, the consequences driving selection involve relative fitness, and copying involves transfer of DNA from parent to child. Trial-and-error learning involves three similar factors: variation in behavior, consequences involving positive and negative rewards, and lasting change in neural structures (modified structures in the individual “descend” from earlier structures).

Variation in behavior occurs for a number of reasons. Many variations are induced by the situation; they may arise from the situation due to innate modules, the current motivational state of the individual, prior

¹⁵ I do not accept the metaphysical possibility that Swampman, at creation, has a mind or that its organs have any functions. Mind presupposes explanatory, non-accidental relations between the individual and a subject matter. Swampman by definition is a cosmic accident. Swampman at creation has no thoughts, beliefs, perceptions, memories, and so on. Similarly functions presuppose explanatory, non-accidental beneficial effects. Swampman by definition bears no explanatory relations to anything at all. At creation there is nothing his heart is supposed to do.

Pavlovian conditioning, or even prior operant conditioning. Another source of variation is the variability inherent in all human behavior. Jerome Frieman reports in his textbook: “Individuals do not perform the same action exactly the same way each time they do it. Even when the individual is well practiced and the stimulus situation is identical on each occasion, there will still be some behavioral variability in how a behavior is performed” (2002: 260). The first source of variability is called *induced* variability. The second is called *behavioral* variability.

Trial-and-error learning involves “trials” – variations in behavior – “errors” – negative reinforcers – and “successes” – positive reinforcers. Trial-and-error learning requires the individual to find behaviors that reduce negative reinforcers and increase positive reinforcers. If the individual can learn, then over time – sometimes very quickly – the individual will find the correct behavior that avoids negative reinforcers and obtains positive reinforcers. Induced variation produces the “trials” that eventually lead to a solution. Behavioral variability then makes more efficient behavior possible through hill-climbing; once the solution is found, behavioral variability produces a more efficient solution. Negative and positive reinforcers select among variants. Successful behaviors are selected by their beneficial consequences; successful behaviors, in virtue of their consequences, are more likely to occur again in similar situations (Frieman 2002: 263–64).

What feedback mechanism makes this possible? What feedback mechanism “integrates information about the behavior” with its consequences? Whatever the details, it involves sensation or perception and memory. The individual must sense or perceive both the behavior and its consequences, and the individual must record and process that information and translate it into future behavior (Frieman 2002: 270). When we learn through trial-and-error learning, we rely on perception and memory to select the right behavior among its variants in virtue of its consequences. Sometimes it is automatic, and sometimes very slow. Sometimes it is entirely conscious, sometimes entirely unconscious, hidden from view. Even single-celled organisms “sense” and “remember” and so learn by trial-and-error.

How does trial-and-error learning generate functions? Take the neural structure underlying my ability to tie my shoelaces, or the behavior (the motion of my fingers) that it causes when I want to tie my laces. (1) It ties my shoes; (2) tying my shoes benefits me (I get what I want, I avoid frustration, I earn the praise of my parents and others, my shoes stay on, I don’t trip, etc.); and (3) I have the structure or can perform the behavior because tying my shoes benefits me. It’s then the function of the structure or the behavior to tie my shoes.

Take any skill you've acquired through learning: passing a soccer ball with your feet or catching a baseball with a glove; pronouncing English verbs and Chinese tones; speaking in public or writing elegant prose; hitting a distant target with bow and arrow; the list is endless. On the present account, nearly every one of these behaviors or the underlying structures will have functions, where the function is often named by the name of the skill: the function of structure underlying my ability to pass a soccer ball is to pass soccer balls. Many of these skills are acquired without consciously and deliberately setting out to acquire them – think of first-language learning, or the learning of various habits that benefit you in one way or another, habits acquired from positive and negative reinforcers that drove selection of the behavior, without your awareness that you were headed in that direction. Other skills are acquired consciously and deliberately. You may very much want to be a good soccer player, an excellent first-baseman, or a world-class archer. Either way, the underlying structure has a function, for the structure results from a feedback mechanism on beneficial effects. Performance normativity is then a species of functional normativity.

Trial-and-error learning, like the normal operation of your metabolism, is a non-hereditary feedback mechanism generating functions; you can acquire all sorts of skills your parents never dreamed of. And now we have another mechanism that doesn't require much history, and so we have another way Swampman can acquire functions; if he has sensory and perceptual capacities, memory, and the mechanism underlying operant conditioning, then he can learn through trial-and-error learning.

Learning also takes us beyond the scope of the “narrowly” biological, where the narrowly biological covers anatomy and physiology, to the “broadly” biological, where biology includes psychology, anthropology, and sociology, especially the learned behavioral traits of individuals. “Nature” has given way to “nurture” as a source of functions.

10 Learning and derived functions

Trial-and-error learning, of course, isn't the only form of learning. Psychologists tend to define learning very broadly. As a result they think there are many forms indeed. In general they see learning as a relatively permanent change in the organism that isn't due to normal development.

You may be wondering about so-called one-off learning, where I learn how to do something without the process of trial-and-error. Is one-off learning a source of functions? Yes it is. Let me explain. The ability to

learn – and to learn in various ways – is itself functional. Though classical empiricists and psychological behaviorists tend to emphasize learning at the expense of the innate, they agree that the ability to learn is built in. But then, we might ask, why is it built in? Mother Nature builds it when the organism needs it to survive. Some organisms really need to learn various things, while others do not. Birds that nest on ledges on cliffs, for example, do not need to learn to recognize their chicks, for the only chicks they'll ever significantly interact with are their own. Birds that nest on crowded beaches, on the other hand, do need to learn to recognize their chicks, for they will see and interact with plenty of chicks that are not their own. "Learning is an option, like camouflage or horns, that nature gives to organisms as needed – when some aspect of the organism's environmental niche is so unpredictable that anticipation of its contingencies cannot be wired in" (Pinker 1997: 242).

But if that is so, then our various abilities to learn will have learning – and so adapting to our environment in beneficial ways – as a function. The general ability to learn, when it leads to learning new abilities, results in those abilities having "derived" functions named, in part, by their beneficial effects. What is learned has a function derived from the general ability. In Millikan's (1984) jargon, learning has adapting to the environment as its *direct* function. The structure or behavior that results has a *derived, indirect* function.

This means "one-off" skills and abilities – even perceptual categories and belief-forming capacities, if there are any – have *derived* functions. It also means that functions from trial-and-error learning will have their functions twice determined. For since trial-and-error learning itself involves consequence selection as a means of adapting to the environment (and so generates direct functions), and since the capacity to learn from trial and error results from consequence selection (and so generates indirect functions), items that result from trial-and-error learning will have their functions over-determined as both direct and indirect functions.

II Conclusion

Teleological views of mind and psychological capacity pervade the history of philosophy. It's thus not unusual to find broadly functional views on the nature of warrant and other epistemological properties embraced in the history of philosophy. Proper function, virtue epistemology has a long history. My view of warrant falls within this teleological frame of mind.

Discussions about the role of functions within epistemology focus on two sources: God and Nature. And when it comes to Nature, natural selection is the paradigm case. And then for the naturalist there is often only the paradigm, natural selection.

But once we review the sources of etiological functions beyond directional natural selection, we see that belief-forming processes may have forming true beliefs reliably as an etiological function from any number of functions. There are many ways belief-forming processes may acquire an etiological function, and so many ways they may acquire the etiological function of forming true beliefs reliably, and so many ways warrant may arise. Warrant requires functions, and functions require history, and so my account of warrant requires history. How much history is required, however, turns on the details of the particular case. I have not examined those details, but I have shown that, at least in principle, the history required may be considerably less than one might have otherwise thought.