

## FUNCTIONS, WARRANT, HISTORY

Peter J. Graham  
University of California, Riverside  
February 2013

Forthcoming in A. Fairweather, ed.  
*Naturalizing Epistemic Virtue*  
Cambridge University Press

I'm a reliabilist about warrant, where warrant constitutively involves the reliability of the belief-forming process.<sup>1</sup> Though popular, reliability theories face (at least) two familiar challenges: the demon-world or brain-in-a-vat case (which apparently shows reliability is not necessary for warrant) and clairvoyant-powers type cases (which apparently show reliability is not sufficient).

I think the way to meet these two challenges is to embed the theory of warrant in a more general theory. Like most epistemologists, I see warrant normatively. And so to meet the challenges I think we should embed the theory of warrant in a more general theory of normativity. Ernest Sosa (2007) embeds epistemic normativity in performance normativity. I see things somewhat differently. I see warrant as a species of functional normativity. I think warrant falls within the more general theory of functions and functional norms.

Here's a quick sketch of my view. Suppose a belief-forming process has forming true beliefs reliably as a function. Given its function, the process will form true beliefs reliably in normal conditions when functioning normally. When forming true beliefs, the process will fulfill its function, and thereby meet a standard or norm constitutively associated with its function. And when functioning normally—whether in normal conditions or not—the belief-forming process will fulfill another standard or norm constitutively associated with forming true beliefs reliably; it will work or operate the way it is supposed to operate, where the right way of working is partly constituted by reliably getting things right in normal conditions. I believe epistemic warrant consists in

---

<sup>1</sup> I use 'warrant' the way most epistemologists use 'justification' or 'justifiedness'. Warrant is then a part of knowledge, but not sufficient for converting belief into knowledge, for warrant may fall short of 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. For more on my use of 'warrant' and related words, see my 'Psychological Capacity and Positive Epistemic Status'.

functioning normally when the process has forming true beliefs reliably as a function. Warranted beliefs are then well-formed beliefs, where being well formed means functioning well, where functioning well is partly constituted by, but not the same thing as, reliably getting things right.

On this view the two challenges are easily met. In the “demon-world” or “brain-in-a-vat” scenario, the subject’s belief-forming process may still function normally, even though the subject is not in normal conditions, and so not forming beliefs reliably true. A “clairvoyant” or “accidentally reliable” belief-forming process has no function, and so nothing counts as functioning normally, and so those beliefs are not warranted, even if reliably true.

All of this, however, turns on the nature of functions. I endorse the *etiological* 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 turns on functions, and functions turn on history, then warrant requires history.

How much history, however, do functions constitutively require? And so how much history does warrant constitutively require? Does warrant require a history of natural selection over many generations? Many of my readers—friend and foe alike—have taken me to suppose it does.<sup>2</sup> But isn’t that asking for too much? How could warrant require a history of natural selection over generations? And 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?

As it turns out, I don’t require natural selection *per se* for functions, and so I don’t require natural selection *per se* for warrant, for natural selection is not the only source of etiological functions. In this paper I discuss two other sources of functions beyond natural selection: an organism’s normal metabolism and trial-and-error learning. Warrant may arise from any of these three sources of etiological functions. Indeed, warrant may arise from any source of etiological functions. If warrant turns on functions, and functions turn on history, then warrant turns on history. How much history? In some cases, not much. Or so I’ll argue.

In the first section I explicate etiological functions. You’ll see why etiological functions require history. In the second and third I show why I think warrant requires functions. In the remaining sections I show why etiological functions, though they require some history, needn’t always require a lot. Even Swampman’s organs, traits, behaviors, and psychological capacities—insofar as he has any—can acquire functions. So fear not: we can understand functions in terms of history, warrant in terms of functions, and so

---

<sup>2</sup> For example, see the friendly presentation of my view in Evans and Smith 2012: 194, and the implicitly critical in Lyons 2011.

warrant in terms of history, without constitutively requiring in each and every case of warranted belief a history of evolution by natural selection.

A caveat. I won't argue that this or that belief-forming process has forming true beliefs reliably as an etiological function. In this respect I'm like any other reliabilist who argues that reliability is required for warrant, but does not go on to argue that this or that process reliably induces true beliefs. I am aware, however, that whether this or that process has the etiological function of being reliable is controversial, and obviously more controversial than the claim that the process is reliable (in part because a process has the function of being reliable only if reliable in normal conditions). I am also aware that the overall persuasiveness of my case depends on showing this or that process has forming true beliefs reliably as a function from this or that source of functions. But I can't do everything at once, and so I won't. Instead I argue that since there are a variety of sources of functions, it is at least possible that this or that source of beliefs has an etiological function from one or more of those sources. And so I hold that warrant, though it requires history, need not require that much.

## I. FUNCTIONS FROM HISTORY

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. 'What is this for, what's it supposed to do, what's its *function*?' The heart is supposed to pump blood; that's its function; that's what it's for.<sup>3</sup>

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. This gets called the function/accident distinction.

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 X's being there, i.e. X's are disposed, do, or can do Z).
- (2) X is there because it does Z (that X's are disposed, do, or can do Z explains why X is there).

---

<sup>3</sup> 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 discuss such views, and how they effect my view of warrant, elsewhere.

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 is there or came to be 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 *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 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, 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. 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. Functions arise through a feedback mechanism that involves explanatorily *beneficial* effects.

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).

Before turning to known feedback mechanisms on beneficial effects, I should review my account of warrant. After all, without seeing why warrant should turn on functions, one might wonder what the fuss is all about.

## II. EPISTEMIC NORMATIVITY AS FUNCTIONAL NORMATIVITY

In this section I argue that three kinds of epistemic normativity are species of functional normativity. First, I explain why normal functioning is constitutively associated with function fulfillment for items with etiological functions. Second, I say why function fulfillment and normal functioning satisfies norms and identify three functional norms. Third, I apply this to belief-forming processes with forming true beliefs reliably as an etiological function. In the section following I explain why I think warrant consists in normal functioning when the belief-forming process has forming true beliefs-reliably as a function.

And so why think normal functioning is constitutively associated with function fulfillment? The answer involves the way normal functioning and normal conditions are fixed on the etiological account. On the account, what counts as normal functioning and normal conditions falls out of the historical explanation for how the item came to have its function. Functions arise when an item produces a beneficial affect 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. 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.

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. Voila, normal functioning and normal conditions (Millikan 1984).<sup>4</sup>

Given the way normal functioning and normal conditions are determined, normal functioning and normal conditions are then constitutively, explanatorily interrelated with function fulfillment. In particular, normal functioning is then 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, normal conditions, and function fulfillment are all holistically interrelated.<sup>5</sup>

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

---

<sup>4</sup> It’s then a priori on this account that *if* an item X has Z as a function, *then* the item fulfilled its function in normal conditions when working normally. For if it did not fulfill its function while working normally in normal conditions it would not have entered into a feedback loop for its functional effect; feedback mechanisms cannot operate on effects that do not occur. It is then a priori that *if* the item has that function, then *ceteris paribus* it will fulfill its function often enough while functioning normally in normal conditions. We can then think of this fact as a formula: *ceteris paribus* normal functioning plus normal conditions leads to function fulfillment, often enough. How often is often enough? This varies. For some items, like the heart, they must fulfill their function nearly all the time when functioning normally in normal conditions; hearts were selected for reliably pumping blood. But for other items, like sperm, often enough is not very often at all. Sperm, or the sperm-producing organ, was not selected for *reliably* fertilizing eggs, even though *fertilizing eggs* is its function.

<sup>5</sup> 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. Thus whatever fixes normal functioning and normal conditions isn’t holistically interrelated with function fulfillment, for conscious, design functions.

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.

I now turn to functional normativity.

In one sense of 'norm', norms are prescriptions or guides. But in another sense, norms are standards or levels of "possible performance that is in some way adequate for fulfillment of a function or purpose" (Burge 2010: 311). 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. When there are functions, there are norms in this sense.<sup>6</sup>

I identify three 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

---

<sup>6</sup> Functional normativity is a more general kind of normativity than the kinds usually studied by philosophers, for philosophers usually study instrumental, moral, and social normativity, kinds necessarily associated with represented goals and intentions, endorsement of standards, and even patterns of approval and disapproval for conformity with the represented norm. Philosophers tend to study prescriptive or guiding norms. Functional norms are a broader kind, associated with functions.

Tyler Burge (2010) 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. Rather he intends to contrast them with other kinds of norms. Positively construed, natural norms are standards or levels of adequacy for performance in fulfilling functions. Negatively construed, natural norms do not depend on the aims or intentions of individuals, on being represented or being endorsed. They apply independently of such aims, intentions, representational capacities, or endorsement. In this way they contrast with norms of instrumental rationality, moral norms and social norms. For instrumental norms depend upon intentions and plans; moral norms (arguably) depend on being represented and understood; and social norms depend upon patterns of approval and disapproval in the relevant population. Furthermore, natural norms do not function as guides. They do not prescribe activity—they are not *dos* and *don'ts*—and they do not guide the entity or individual they apply to. No norm prescribes pumping blood to the heart, and the heart does not represent, internalize, or consult any norms or prescriptions. Natural norms need not be represented or positively endorsed by the entity or individual they apply to. Some entities like hearts cannot even represent or endorse such norms, for they cannot represent or endorse anything at all.

Once we fully appreciate the possibility of this broader category of normativity, one independent of any individual's goals, endorsement, or capacity to represent such norms, we can appreciate what Burge considers a "momentous structural feature" of the world: wherever there is natural teleology—function or purposiveness—there are "standards for realizing" functions: natural norms.

*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.

And we understand the second norm in terms of the first—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. Normal functioning *encodes*, for partly constituted by, function fulfillment, for items with etiological functions.

I now explain why kinds of epistemic normativity are species of functional normativity.

Epistemic norms are norms understood in terms of promoting true belief and avoiding error. Many epistemic norms are prescriptive. Some prescribe inquiry; they 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 for and against. Some prescribe thresholds for decision, conviction or judgment. Many of these norms not only prescribe but 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. Though many norms prescribe and guide, some do not. Some norms set standards or levels of adequacy for performance in fulfilling functions. Some norms are functional norms. Could there be epistemic norms that are norms in this sense?

Assume a belief-forming process has forming true beliefs reliably as an etiological function. On this assumption there are three natural norms governing the process: function fulfillment (and so reliably forming and sustaining true beliefs), normal functioning, and function fulfillment *because* normal functioning. And since normal functioning for items with etiological functions is associated with and understood in terms of function fulfillment, normal functioning for the belief-forming processes is associated and understood in terms of reliably inducing true beliefs; normal functioning *encodes* reliably getting things right. 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.

### III. WARRANT FROM FUNCTIONS

Why think, as I do, that warrant consists in normal functioning when the belief-forming process has forming true beliefs reliably as an etiological function?



I believe there are two general (and typically opposing) characterizations or conceptions of warrant within epistemology. There are epistemologists who think warrant is constitutively tied to reliably promoting truth, to getting things right. Like consequentialists in ethics, they believe epistemic “rightness” turns on promoting “epistemically good” consequences. And then there are epistemologists who think warrant is more strongly tied to following “epistemically good” or “correct” procedures of belief-formation and maintenance. Like deontologists in ethics, these epistemologists believe epistemic “rightness” turns on following or conforming to “right rules” or “rational procedures” even if, as a matter of fact, following those procedures fails to produce reliably true beliefs.

Tyler Burge (2010: 50-1) draws the contrast as follows. On the one hand, there are those who see warrant “vertically” as consisting in certain kinds of mind to world relations; warrant consists in certain relations between belief and the subject-matter those beliefs are about. On the other hand, there are those that see warrant “horizontally” as consisting in certain kinds of mind to mind transitions or in certain mind to mind relations; warrant consists in certain transitions from one mental state to another.

Consider perceptual warrant. Typically perceptual beliefs are caused and sustained by perceptual representations. Why should perceptual representations, and the transition from representation to belief, warrant perceptual beliefs? Reliabilists point to the reliability of perceptual representations and the reliability of beliefs so caused and sustained. Proceduralists point to the relation between perceptual representation and perceptual belief and the transition from the representation to the belief.

The demon-world and clairvoyant-powers cases are used to illustrate the difference between the two conceptions and to support the latter, proceduralist conception. In the demon-world case, the individual relies on correct procedures of belief formation, but gets everything wrong. How then could warrant constitutively require reliably getting things right? In the clairvoyant-powers case, the individual does not rely on any recognizably correct procedure, but as it turns out reliably gets things right. How then could sheer reliability ensure warrant? These questions lead many to conclude that warrant and truth are only loosely related at best. John Pollock and Joseph Cruz even conclude that getting things right ‘is not a very interesting part of our best understanding’ of warrant (Cruz and Pollock 2004: 137), and even that warrant bears ‘no connection to the truth’. They conclude instead that warrant turns on rational procedures, with the psychological processes on which beliefs are formed and sustained (Pollock 1999: 385-6; Pollock and Cruz 1999: 14-15).

Those in the first camp find the idea that warrant has no connection with truth deeply puzzling. After all, if warrant bears no connection with truth, what makes the warrant *epistemic*? If warrant bears no connection with truth, how could warranting states and events *epistemically* support belief? Those in the first camp insist that warrant

*connects* with truth, that warrant somehow or other *aims* at, *contributes* to, or *conduces* true belief.

In our history various philosophers have tried to ensure a connection between the two conceptions. Descartes saw warrant in terms of conforming to correct procedures—forming belief on the basis of clear and distinct perceptions. He also saw warrant as tied to truth and reliability, for correct procedures were sure to reliably lead to truth because of God’s goodness. For Descartes, God ensures a fit between the two conceptions. Kant and other idealist minded philosophers also saw a connection between the two conceptions. For Kant, the objects of cognition—the subject matter of thought—was importantly constituted by patterns of good cognition. So insofar as warrant consists in thought’s conformity to correct procedures, our thoughts are sure to be largely correct about their subject matter. For Kant, transcendental idealism ensures a fit between the two conceptions.

I see the urge in the history of philosophy to connect or even unify our two conceptions as reason for preferring a unified account against alternatives. I believe a unified, non-theistic non-idealist account is possible. I believe we can unify our two conceptions if we construe warrant as consisting in the normal functioning of the belief-forming process when the process has forming true beliefs reliably as an etiological function. *Epistemically* correct procedures are then just those mental or psychological elements in the normal functioning of a belief-forming process when normal functioning encodes reliably getting things right. Normal “procedures” are epistemically correct procedures when the process has forming true beliefs reliably as a function, for those procedures are constitutively associated with reliably getting things right. For processes with reliably getting things right as an etiological function, normal functioning just is working or operating the way the process worked or operated so as to produce true beliefs in normal conditions.

Whereas Descartes and Kant first identified correct procedures and then explained why following correct procedures was connected with truth and reliability, on the view I favor we first connect truth and reliability with functions and then explain why the procedures are correct in terms of their constitutive and explanatory connection to truth and reliability in normal conditions. Instead of God insuring the reliability of good procedures, or the good procedures of the mind constructing its subject-matter, an explanatory history of reliable success constructs correct procedures.

This view of warrant easily explains why warrant aims at, contributes to, and conduces truth. Consider “aiming.” Normal functioning for any item with an etiological function “aims” at function fulfillment. So if the function is reliably getting things right, normal functioning “aims” at reliably getting things right, and so “aims” at truth.

Consider “contributing.” Normal functioning for any item with an etiological function explanatorily and non-accidentally contributes to function fulfillment. Normal functioning is the item’s contribution to function fulfillment; the environment contributes

the other part. So if the function is to reliably get things right, normal functioning non-accidentally contributes to reliably getting things right.

Consider “conducting.” Normal functioning for any item with an etiological function explanatorily and non-accidentally conduces function fulfillment in normal conditions. Normal functioning is the item’s non-accidental, explanatorily relevant contribution to function fulfillment, and thereby non-accidentally conduces function fulfillment in normal conditions. So if the function is to reliably get things right, normal functioning non-accidentally conduces towards getting things reliably right, and so conduces true belief, in normal conditions.

This view of warrant also easily explains why warrant may persist outside of normal conditions. For on this view, though warrant is constitutively associated with reliably getting things right, warrant is not co-extensive with reliability. The “demon-world” or massively deceived “brain-in-a-vat” case shows that reliability in present circumstances—contextual or *de facto* reliability—is not necessary for warrant. But if warrant consists in normal functioning, warrant does not automatically lapse when envatted or massively deceived due to unusual circumstance, for a process may still function normally even when the system is no longer in normal conditions. Recall the heart removed from the chest during the complicated surgery; it may operate normally, even if not in normal conditions. So if a belief-forming process has forming true beliefs reliably as a function, then even envatted or massively deceived it may still confer warrant when functioning normally.

BonJour-type “clairvoyance” cases also challenge simple, *de facto* reliability theories. They seem to show that *de facto* reliability is not sufficient for warrant.<sup>7</sup> If warrant consists in normal functioning, then warrant entails that the belief-forming process have a function. But a process could reliably induce true beliefs without having that function, or any function at all. And then nothing would count as normal functioning. And that’s one thing I think is going on in BonJour-type cases: the process arises by accident and doesn’t exist or continue to exist because it reliably induces true beliefs, and especially not because they confer some good or benefit on the individual who has the process. If warrant consists in normal functioning, such processes, no matter how reliable, do not confer warrant, for they have no function.<sup>8</sup>

And so on this view we can have our cake and eat it too. Both conceptions of warrant are correct: warrant is constitutively associated with reliably getting things right, and warrant consists in following epistemically correct procedures, for correct procedures are constitutively associated with reliably getting things right.

---

<sup>7</sup> BonJour also thought they showed that warrant also requires the individual to have reflective access to internal mental states that could serve as justifying grounds the individual could cite or appreciate in favor of the truth of her beliefs. I shall table BonJour’s more ambitious intention and simply address the point that reliability doesn’t seem sufficient for warrant.

<sup>8</sup> I discuss these cases in more detail in my 2011a, b, c and my 2012.

So that's why I think warrant consists in the normal functioning of the belief-forming process when the process has forming true beliefs reliably as an etiological function. And 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 functions require. In the remainder of the paper 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.

#### IV. NATURAL SELECTION

Natural selection requires three elements: variation, heredity, and fitness (beneficial) consequences. 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 hereditary, 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 *heritable*; 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, unlike directional selection, is recent, ongoing selection. Your heart hasn't simply been selected by nature eons ago; it's been maintained in your

very lifetime. Human hearts, visual systems and reasoning capacities are undergoing maintenance selection as we speak.

Natural selection of both kinds generates functions because it explains why a trait with a certain capacity exists or continues to exist in a population because of its beneficial features or effects on continued life and the propagation of life; natural selection is a feedback mechanism on beneficial consequences. Applied to our beetles, the function of green coloration is to provide camouflage, because (1) being green provides camouflage; (2) providing camouflage benefits the beetles; and (3) they are green and thereby camouflaged because providing camouflage benefits the beetles. A belief-forming psychological capacity will then have the function of reliably inducing true beliefs via natural selection 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.

## V. SELF-REPLICATION AND REPAIR

In now discuss another source of etiological functions. In *What Functions Explain*, McLaughlin argues for a non-hereditary feedback mechanism alongside natural selection (2001: 162-190). 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 comprised of one group of cells as input and produces later cycles comprised of 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.<sup>9</sup>

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 only benefits 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.

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

---

<sup>9</sup> 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.

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 only works 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, McLaughlin 2001: 67-8, 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.<sup>10</sup>

---

<sup>10</sup> I say more about Swampman in my 2012.

## VI. 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 heredity: the variation is genetic, the consequences driving selection involve relative fitness, and heredity 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 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: 263-4).

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.), (3) and 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; 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 behavior or structure underlying passing a soccer ball with your feet 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 resulting behavior or 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. A belief-forming psychological capacity will then have the function of reliably inducing true beliefs via trial-and-error learning provided (1) it reliably induces true

beliefs; (2) reliably inducing true beliefs benefits the individual; and (3) the individual possesses such a capacity because reliably inducing true beliefs benefits the individual.

## VII. 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 1994: 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.

I conclude 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.

## VIII. CONCLUSION

Teleological view 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. My view of warrant falls within this teleological frame of mind.

Contemporary debates about functions focus on two possible 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. No wonder the natural response to my account of warrant is to read it as requiring natural selection over many generations. And I am sure I only entrench this reading of my view by focusing on the heart and its function as my paradigm case.

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. My account of 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.<sup>11</sup>

---

<sup>11</sup> I presented an earlier version 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.

## REFERENCES

- Bedau, Mark (1991). 'Can Biological Teleology be Naturalized?' *Journal of Philosophy* 88: 647-655.
- Burge, Tyler (2010). *Origins of Objectivity* (New York: Oxford University Press).
- Cruz, Joseph & John Pollock (2004). 'The Chimerical Appeal of Epistemic Externalism' in R. Schantz, ed., *The Externalist Challenge* (Berlin: De Gruyter).
- Evans, Ian and Nicholas Smith (2012). *Knowledge (Polity)*.
- Frieman, Jerome (2001). *Learning and Adaptive Behavior* (Wadsworth Publishing).
- Graham, Peter J. (2010). 'Testimonial Entitlement and the Function of Comprehension' in A. Haddock, A. Millar, D. Pritchard, eds., *Social Epistemology* (Oxford: Oxford University Press).
- Graham, Peter J. (2011b). 'Psychological Capacity and Positive Epistemic Status' in J. Graper Hernandez, ed., *The New Intuitionism* (Continuum).
- Graham, Peter J. (2011c). 'Does Justification Aim at Truth?' *Canadian Journal of Philosophy* 41: 51-72.
- Graham, Peter J. (2011d). 'Perceptual Entitlement and Basic Beliefs' *Philosophical Studies* 153: 467-475.
- Graham, Peter J. (2012). 'Epistemic Entitlement' *Nous* 46: 449-482.
- Graham, Peter J. (forthcoming-a). 'Perceptual Entitlement and Natural Norms' in Peter Graham and Nikolaj Jang Pedersen, *Warrant Without Reason* (Oxford: Oxford University Press).
- Graham, Peter J. (forthcoming-b). 'The Function of Perception' in Abrol Fairweather, *Virtue Scientia*.
- Lyons, Jack (2011). 'Response to Critics' *Philosophical Studies* 153: 477-488.
- McLaughlin, Peter (2001). *What Functions Explain*. (Cambridge: Cambridge University Press).
- Millikan, Ruth (1984). *Language, Thought, and Other Biological Categories*. (Cambridge, MA: The MIT Press).
- Pinker, Steve (1994). *How the Mind Works* (Norton).
- Pollock, John (1999). 'Procedural Epistemology' in J. Greco & E. Sosa, eds., *The Blackwell Guide to Epistemology* (Blackwell).
- Pollock, John and Joseph Cruz (1999). *Contemporary Theories of Knowledge*, 2<sup>nd</sup> (Rowman & Littlefield).
- Ruse, Michael (1971). 'Functional Statements in Biology' *Philosophy of Science* 38: 87-95.
- Sosa, Ernest (2007). *A Virtue Epistemology* (Oxford: Oxford University Press).
- Wright, Larry (1973). 'Functions' *The Philosophical Review* 82: 139-168.
- Wright, Larry (2012). 'Revisiting Teleological Explanations: Reflections Three Decades On' in P. Huneman, ed., *Functions: Selection and Mechanisms* (Synthese Library).