




8

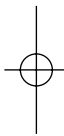
Testimonial Entitlement and the Function of Comprehension

Peter J. Graham




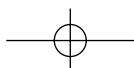
Why do you think it will rain tomorrow? Because, you say, it always rains this time of year; you can see dark clouds for miles; and besides, you read the weather report. The reasons you offer *justify* your belief; they comprise a *justification* for your prediction about the weather. Justifications are commonplace. Justifications involve reasons or evidence you often cite when asked how you know. Justifications support your ability to *justify*—to argue in favour of—your belief, to show that your belief is (likely to be) true (Unger 1968; Audi 1988; Kornblith 2008).

Justification is one kind of positive epistemic status or standing. Positive epistemic standings are goods, successes, fulfilments, or achievements understood in terms of promoting true belief and avoiding error (Alston 1985, 2005).



Epistemic justification is not the only kind of positive epistemic standing. Perceptual beliefs paradigmatically enjoy positive epistemic status, but higher non-human animals, human infants, children, and ordinary adults often lack evidence or reasons in support of their perceptual beliefs, and even more often, if not always, lack the ability to cite evidence or reasons in support of their perceptual beliefs. Ask a small child why he believes the ball continues to exist as it passes behind the screen, and at best you'll get a gurgle or a burp. Ask a chimpanzee trying to open a crate why she believes there are bananas inside, and you won't get any answer at all. Well-formed, garden-variety perceptual beliefs enjoy positive epistemic status, but that status does not depend on reasons or evidence, and it especially does not depend on the subject's ability to offer reasons or evidence (Goldman 1979).

I use the term 'warrant' for positive epistemic statuses that provide epistemic support or grounding in favour of holding the belief. Justifications are a familiar kind. Evidentialism arguably captures the contours of this kind of warrant. But justification is clearly not the only kind of warrant. I use 'entitlement' for positive epistemic support or grounds in favour of belief when the support or grounding



does not involve or depend upon evidence or reasons, especially evidence or reasons the subject might cite in favour of his or her belief (cf. Burge 1993, 1996). Entitlements are goods, successes, or fulfilments in the way of promoting true belief and avoiding error that do not depend on evidence or reasons available to the subject. Though many perceptual beliefs enjoy justifications, paradigmatically they enjoy entitlements.

What about so-called ‘testimony-based beliefs?’ These are beliefs based on our capacity to comprehend assertive speech acts. What kind of warrant do they enjoy? When you tell me the rain won’t let up and I comprehend and accept your assertion, what kind of warrant does my comprehension-based belief enjoy, justification or entitlement?¹

On the traditional view, warrants for comprehension-based beliefs rely exclusively on entirely first-hand justifications. This tradition has fallen on hard times. For though comprehension-based beliefs are by and large warranted, it is now widely thought that fully adequate, entirely first-hand justifications are in short supply (Coady 1992; Fricker 1995).

Though comprehension-based beliefs sometimes enjoy fully adequate, entirely first-hand justifications, and very often enjoy partially adequate, partially first-hand justifications (cf. Faulkner 2002), I shall argue that comprehension-based beliefs formed though a process of filtering generally enjoy *prima facie pro tanto* entitlement.² We thus need not worry over the paucity of fully adequate, entirely first-hand justifications. Though justifications clearly play an important role in a complete account of the epistemology of comprehension-based beliefs, they do not play the *only* role.

Why should beliefs based on comprehension-via-filtering enjoy entitlement? I shall first argue for an account of (a kind of) entitlement and then apply it to comprehension-based beliefs. A belief enjoys (a kind of) *prima facie pro tanto* entitlement when based on a normally functioning belief-forming process that has

¹ Positive epistemic statuses make up the subject matter of epistemology. Some apply to beliefs, some to believers. Knowledge and truth are both positive epistemic statuses that apply to belief. Warrant is another. The connections between knowledge, true belief, entitlement, and justification are complex. Justifications and entitlements both fall short of enabling knowledge, even when they warrant a true belief. Such is the stuff of Gettier cases. Justification (as understood here) is clearly not necessary for perceptual knowledge (Kornblith 2008). I can only explicate testimonial entitlement here, and cannot pursue its connections with testimonial justification and testimonial knowledge. I have discussed testimonial knowledge elsewhere (Graham 2000a, 2000b, 2006b). See note 4 for my construal of the adjective ‘testimonial’.

² ‘*Prima facie*’ entitlement means defeasible entitlement. Counter-considerations may overturn the entitlement. Justification or further entitlements may be required to restore the warranting force of an entitlement. ‘*Pro tanto*’ entitlement means *some* entitlement, where some may fall short of enough for on balance entitlement, given some threshold (Graham 2006a). Justification or further entitlement may supplement the entitlement, converting some warrant into enough. Do not confuse my thesis that comprehension-based beliefs generally enjoy *prima facie pro tanto* entitlement with the obviously absurd thesis that comprehension-based beliefs *ipso facto* enjoy indefeasible on balance warrant, that believing everything you hear somehow confers maximal epistemic warrant.

forming and sustaining true beliefs reliably as a function.³ And comprehension-with-filtering has that function. Beliefs based on comprehension-with-filtering thus enjoy (a kind of) entitlement when our capacity to comprehend and filter functions normally, for our capacity has forming and sustaining true beliefs reliably as a function. When functioning normally, comprehension-with-filtering confers *prima facie pro tanto* entitlement on the beliefs it causes and sustains.

In the first section I say a few words about our psychological capacity to comprehend and filter. In the second I explicate function, normal conditions, and normal functioning. In the third I put these notions to use to argue for my first main thesis, that (a kind of) epistemic entitlement derives from the normal functioning of the belief-forming process when the process has forming true beliefs reliably as a function. In the fourth I argue indirectly for my second main thesis by first arguing that *assertion* has causing and sustaining *true* hearer beliefs as a function. In the fifth I defend this claim against four counter-arguments. I then argue directly in the sixth for my second thesis that *comprehension-with-filtering* has causing and sustaining true beliefs *reliably* as a function. Though I have tried to be concise, the paper is not short enough to be easily read without taking a break. A natural place to pause would be at the end of section three. Taking a moment to digest the first main thesis may help prepare the palette for the second. Taking a break after the fifth section might not be a bad idea either.

1. COMPREHENSION-WITH-FILTERING

For the most part, we automatically, quickly, and reliably comprehend speech. Comprehension involves a number of factors: standing abilities to comprehend shared meanings of words, shared awareness of relevant features of the context of utterance, and standing expectations of conformity to cooperative principles or maxims governing speech behaviour, not to mention the rest of the language suite: syntax, morphology, phonology, and speech perception.

³ I say 'a kind of entitlement' because I believe one kind derives from normal functioning (as I argue below) and another related kind may derive from function fulfilment. Still a third related kind may derive from function fulfilment due to normal functioning. This third kind is related to what Duncan Pritchard would call an achievement, a success through the exercise of ability. On a looser understanding of 'achievement' where achievements are just successes or fulfilments, both the entitlement that turns on normal functioning and the entitlement that turns on function fulfilment would be achievements. I hope to pursue the connection between kinds of entitlement and notions of achievement elsewhere.

Michael Bergmann (2006) (following Plantinga 1993), Tyler Burge (2003), and Ernest Sosa (2007) each develop accounts of entitlement structurally similar to mine. My account explicitly relies on the etiological notion of function and normal functioning. Plantinga and Sosa explicitly reject reliance on etiological functions. Burge remains neutral on etiological functions and on the importance of normal functioning. I intend to discuss their views in detail elsewhere.

Usually we correctly comprehend what people say. But sometimes we make mistakes: we sometimes miscomprehend and thereby misrepresent the content or force of a speech act, or both.

Comprehension states are exercises of our capacity to comprehend. Since my focus is assertive communication, I shall use ‘comprehension state’ throughout for comprehension states that purport to represent assertions.

I use ‘assertion’ broadly to cover the whole range of constative speech acts: affirm, allege, avow, certify, concur, declare, describe, disclose, inform, predict, recount, report, say, state, submit, tell, and testify, among others.⁴

Comprehension-based beliefs are beliefs caused or sustained by taking comprehension states at face value. In the ordinary case, a speaker (sender) asserts that *P*. The hearer (receiver) represents *P* as asserted.⁵ The representation of the speech act—the hearer’s comprehension state—disposes the hearer to believe that *P*. In the default case, comprehension prompts belief in the propositional content of the assertion as represented. Being so moved constitutes taking comprehension at face value. My topic is the epistemology of beliefs so formed. I thus use ‘testimonial entitlement’ to refer to the entitlement that beliefs based on comprehension enjoy.

⁴ My subject matter is the epistemology of beliefs based on our capacity to comprehend assertive speech acts. This category is much broader than the category of testimony-based beliefs, insofar as (a) we take the latter category to involve a narrow account of testimony (as a speech act), and (b) we require that a testimony-based belief that *P* be caused or sustained by a speaker’s testimony that *P*. Many comprehension-based beliefs do have, as their distal causes, speakers having testified that such and such is the case. But many do not. For testifying (taken narrowly) is just one of many ways a speaker can assertively present a proposition (Alston 2000; Bach and Harnish 1979); testifying is a particular way of assertively presenting a proposition, involving additional intentions on the part of the speaker (Graham 1997). So there are plenty of comprehension-based beliefs that are not ‘based’ on a speaker’s testimony. Furthermore, beliefs based on miscomprehending a speech act are still comprehension-based. A speaker may question whether *P*, but the hearer miscomprehends the force, and takes the question to be an assertion, and comes to believe that *P*. Or a speaker may assert *Q*, but the hearer miscomprehends the content, and comes to believe that *P*. Since both beliefs are exercises of the hearer’s general capacity to comprehend speech, they are both comprehension-based beliefs. The category of comprehension-based beliefs is thus larger than the category of testimony-based beliefs, and larger than the category of assertion-based beliefs. The epistemology of ‘testimony’ literature is, for the most part, interested in this broader category. ‘Testimonial entitlement’ is thus entitlement for beliefs based on comprehension.

⁵ Representing an utterance as the assertion that *P* sometimes leads a hearer to form the belief that the speaker believes *P*, or to believe that the speaker asserted that *P*. But the central case of a comprehension-based belief is a belief in the very propositional content taken as asserted. For discussion of this point, see Hunter (1998).

I allow for leeway as to what counts as the propositional content taken as asserted. It may include simple logical entailments or obvious evidential consequences from what was strictly speaking said or meant. If someone says John was at the party, you may take the speaker to have also asserted that someone was there (Goldberg 2001). It may include obvious conversational implicatures. Someone says there is a gas station around the corner, and you take the speaker to assert that you can fill your tank there. It may also include what was communicated non-verbally, given an already existing linguistic context. Pointing in response to a query is a perfectly good way of telling someone where something is, or where to go. Assertions are also made both directly and indirectly. You can assert by asking a question or making a command. In such a case, you make two speech acts with one utterance.

Our developmental starting point is not a broadly Humean, sceptical stance, where we withhold belief until sufficiently strong supporting reasons come along. Rather our developmental starting point is a broadly Reidian, credulous stance (Harris 2002; Koenig and Harris 2005). Very young children nearly automatically accept whatever they comprehend. We are built to take comprehension states at face value; credulity is the default.⁶

That said, we obviously do not believe everything we hear or read. Filtering soon develops. If you label something right before a child's eyes when she already confidently believes the opposite, she is less apt to believe you. If you call a cup a 'dog' right before a 16-month-old, the child will direct more attention to you than she will to speakers who correctly label the object. If you do the same thing before an 18-month-old, she may even correct you (Koenig, Clement, and Harris 2004).

Another filter involves manner of presentation. Three- and four-year-olds are sensitive to how confident an assertion is. An unsure, less committal tone of voice sometimes makes acceptance less likely. Conversely, more confident assertions make acceptance more likely (Moore, Bryant, and Furrow 1989; Jaswal 2004).

Young children acquire track record information, either first hand or from others. On the basis of this information, they tend to selectively accept assertions. When children have information that a speaker is reliable, that speaker tends to prompt belief more often than the speakers which children 'know' are less reliable. Again, it is three- and four-year-olds that display this ability, with four-year-olds showing more sensitivity to such information (Koenig, Clement, and Harris 2004). This stands to reason, for only after children acquire the concepts of belief and truth are they in a position to form beliefs about which of their interlocutors are reliable, beliefs they can then rely upon in blocking or supporting the transition from comprehension to belief. (I have in mind the false belief task.)⁷ Coherence-checking, a fourth form of filtering, emerges later. I will have more to say about coherence checking further along. As we mature, we rely on these filters and more.

⁶ For a survey of impressive psychological evidence in favour of the claim that credulity is the default, see Gilbert (1991). For additional evidence and argumentation, see Gilbert (1993). Gilbert examines three models of the transition from comprehension to belief: Cartesian, Spinozan, and Cartozan. The Cartesian sees comprehension as neutral towards belief. The Spinozan sees comprehension as consisting in belief; to comprehend an assertion as an assertion that *P* just is to believe *P* as a result of one's linguistic ability to comprehend. The Cartozan sees comprehension states as disposing belief, but still ontologically distinct from belief. He favours the Spinozan view. The research he surveys, and his own research conducted with colleagues, clearly undermines the Cartesian view. But it is not obvious that it favours the Spinozan over the Cartozan position. I favour the less controversial, Cartozan view. I discussed Gilbert's view in Graham (1999).

⁷ Witness Origi and Sperber: 'There are . . . reasons to expect the ability to attribute false beliefs to others to develop after the ability to communicate verbally. The attribution of false beliefs to others plays an obvious role in the ability to filter false information communicated either by mistaken or deceitful speakers' (2000: 163)

Filtering need not involve explicit awareness—belief or judgement—that counter-considerations are absent. Filtering need not involve reasons and reasoning. The subject need only respond by dampening, if not entirely suspending, the normal force of comprehension (Goldberg and Henderson 2006).

In sum, filtering involves sensitivity to counter-considerations: were there are counter-considerations of a certain sort, acceptance would be less likely. Different filters involve different sensitivities. We obviously do not believe whatever we are told; credulity as default does not imply unalterable gullibility. We're no fools.

Sometimes filtering blocks the normal force of comprehension, but we still come to believe what we were told. Consider Nick and Oscar. Nick correctly comprehends Oscar's assertion. But Nick is suspicious of Oscar's credibility. Perhaps Nick antecedently doubts that Oscar will tell the truth. Perhaps Oscar's story does not fit with other things Nick believes. Perhaps Oscar's report contains internal contradictions. Or perhaps Oscar is alarmingly fidgety. If this were the end of the matter, Nick would not believe what Oscar just said. But on reflection Nick thinks Oscar *can* be relied upon. Taking as a premise that Oscar asserted that *P*, Nick comes to believe that *P*. Nick's warrant is a justification, an argument from Oscar's assertion that *P* as premise (and other believed propositions) to the conclusion that *P*. Nick's belief is, as it were, *from* comprehension without being *through* comprehension (cf. Audi 1997).

Once we mature, comprehension involves a 'default-trigger-and-evaluate' structure. Peter Lipton explains:

[I]n most contexts the hearer simply accepts what is [taken as] said, without engaging in any conscious evaluation or inference . . . [But] there are diverse conditions that may trigger the hearer to switch into evaluative mode, where he pauses to consider whether he should believe what he has been told. . . . The obvious triggers include cases where what is claimed obviously contradicts firmly held beliefs, cases of contradictory testimony, and cases where there is reason to worry about incompetence or insincerity. (2007: 241)

Comprehension states dispose belief, but not willy-nilly. Though credulous, we do not believe everything we are told. We filter and, on occasion, evaluate.

2. ETIOLOGICAL FUNCTIONS

I now turn from psychology to teleology. In this section I explicate the notions of function, normal functioning, and normal conditions. There are at least three kinds of functions. My account of entitlement employs the third.

First, there is the familiar category of intended (consciously assigned) functions for items we design to produce certain effects. Artefacts (tools) have many of their functions because of what we have in mind when making them. Usually an artefact's function is what it was consciously and intentionally designed to do. Belts were designed to hold up trousers. Watches were designed to tell the

time. Cars were built to take us from point A to point B. The artefact was made, shaped, arranged, composed, manufactured, in order to produce a certain effect. The effect is what the designer had in mind that explains why she made, shaped, arranged, composed, or manufactured the artefact. Call these d-functions.

Second, there are functions fixed by one-off, occasion specific intentions. Intentions are goals, aims, or purposes. Like function, intention is a teleological notion. An intention on an occasion explains why someone does something, or what he or she is using an item for. I went out to get a breath of fresh air. I put the stapler in the doorway to keep the door ajar. Call these functions i-functions.

A third kind of function applies paradigmatically to biological organs and systems, though it has a much wider application. It even applies to some functions of artefacts; not all functions of artefacts are d-functions. I have in mind etiological functions, where the function of an item essentially depends on certain explanatory features of its history. Call etiological functions e-functions.⁸

An e-function of an item is the effect of ancestors of the item that explains why the item was replicated, and so why the item exists now, where the item (or the system of which it is a part [Buller 1998, 2002]), underwent some form of selection (natural, artificial, cultural, or some other form). Take the heart. Its function is to pump blood. Hearts now exist because ancestors of hearts pumped blood, and in so doing contributed to survival long enough for new hearts. The eye is for seeing. Eyes now exist because ancestors provided information about the spatial array of objects in the environment. For most animals, sexual reproduction is the mechanism of copying where later hearts are produced on the model of past hearts. But other methods of copying will do. E-functions arise when certain effects of ancestors of an item explain why the item persists.⁹

E-functions are real features of biological kinds. When evolutionary biologists taxonomize according to function, they look for effects that played an important role in explaining why the organs, systems, or activities of animals persist.

But as I have already said, e-functions are not simply biological functions. E-functions are effects of ancestors of any item that played a role in explaining why reproductions of that item exist, no matter what the item is, or

⁸ There is the use of 'function' to describe how something works or operates—what it does—as a part of a larger system (Cummins 1975). Functions, in this sense, are the causal role capacities of parts that contribute to some capacity of the containing system. If the part lacks the capacity, it lacks the function. A 'functional analysis' of the system reveals 'functions'—capacities or dispositions—of the parts; the analysis reveals how the system operates or 'functions'. Call these causal-role functions, or just c-functions.

⁹ The etiological account was inspired by Larry Wright's (1973) article 'Functions' and then developed independently along different lines by Millikan (1984, 2002), Karen Neander (1991a, 1991b), Peter Godfrey-Smith (1994), and David Buller (1998, 2002), among others. For a useful anthology, including challenging criticisms of the etiological view, see Buller (1999).

how it was reproduced. Learned behaviours, linguistic devices, conventions, cultural practices, games, and so on, may all have etiological functions (Millikan 1984).

The logic of e-functions comes with an informative account of what counts as normal conditions and normal functioning (Millikan 1984: 33–4). When an item has an etiological function, there is an explanation in terms of that function for why the item exists. It exists now, in part, because its ancestors, when working or operating a certain way in certain circumstances, produced the effect that contributed to the item being replicated. That effect becomes the item's etiological function. The way the item operated according to that explanation counts as *normal functioning*. If the current item works or operates that way (or in relevantly similar ways; there are borderline cases), then it functions normally; it operates the way it is supposed to operate. And the conditions according to that explanation (or conditions of similar type; there are borderline cases), fix *normal conditions*. Take the heart. It beat (worked, functioned) a particular way in an organism's body in a particular environment. By beating that way, it pumped blood and thereby contributed to the survival and reproduction of the organism. Pumping blood thereby became the function of the heart. Normal functioning is then just beating or working the way it did that contributes to the explanation of the heart's acquiring the function of pumping blood.¹⁰

This explanatory, individuating interconnection (between function ascription, function fulfilment, normal conditions, and normal functioning) means that for any item with an e-function, *ceteris paribus* it will fulfil its function, often enough, in normal conditions when functioning (operating, working) normally. If it did not, then the item would not have been replicated because of that effect; it would not have that e-function. This falls out of the logic of e-functions.

¹⁰ There are (at least) four attractive features of the etiological account of functions. First, as we have just seen, it gives an informative and non-circular account of normal conditions and normal functions. Second, it underwrites teleological explanations; something is there because of its function. Third, it accounts for the difference between the function of an item and side effects that may be beneficial for the item or the user of the item. The function of my nose involves affecting incoming air, but not to hold up my glasses. The function of a large belt buckle in the Old West is to keep up trousers, and maybe also to display social status, but not to stop bullets during a shoot-out at the corral. Fourth, it explains malfunction. A diseased or deformed heart still has the function of pumping blood even if it cannot, for the heart is there because its ancestors pumped blood. A broken belt buckle is still supposed to hold up trousers, for it was reproduced on the model of buckles that did.

Another account of functions aimed at accounting for biological functions is the propensity account (Bigelow and Pargetter 1987). On this account the function of an item is its propensity to contribute to the fitness of the organism in its natural habitat. There are four problems with this account. First, the view helps itself to an unexplained account of natural habit. Second, it fails to support teleological explanations, for propensities are future effects, and future effects cannot explain the current existence of anything. Third, it does not capture the function/side-effect distinction: holding up glasses may have a propensity to contribute to fitness, but that is not a function of my nose. Fourth, malformed hearts have no propensity to contribute to fitness, so they have no function on this view (cf. Neander 1991a, 1991b).

This fact marks an important difference between e-functions, on the one hand, and d-functions and i-functions, on the other. You can design something to produce a certain effect, but it may fail to ever do it. Think of the dustbin of failed inventions. You can also intend to do something but not succeed. If only life were so easy. Not so for etiological functions. No type of item with an etiological function can have tokens that fail across the board.

We now have an important explanatory fact to rely upon for items with etiological functions: *ceteris paribus* normal functioning in normal conditions leads to function fulfilment. We can then partially *individuate* and *explain* normal functioning in terms of normal conditions and function fulfilment. If you know the etiological function of an item, functioning normally *just is* working or operating in the way that, according to the historical, normal explanation for the item, fulfils that function often enough *ceteris paribus* in normal conditions. Normal functioning is thus individuated and understood in terms of function fulfilment.

You can fulfil your function without functioning normally, and function normally without fulfilling your function. Stimulate a diseased heart artificially and it may pump blood, though not because it is working the way it is supposed to. You can also take a heart out of someone's chest and put it in a sterile dish; it may still function (work, operate) normally for a time, but no blood is passing through it. Put a car up on a lift for inspection, and it may function just fine, even if it is not hurtling you down the highway at 65 miles per hour. Fulfilling a function is one thing, functioning normally is another, despite their explanatory relation.

Different items fulfil their functions at different rates, even in normal conditions when functioning normally. The function of sperm—or the sperm-producing device—is to fertilize eggs.¹¹ But sperm hardly ever do that, even in normal conditions, functioning as normally as can be. Sperm need only fertilize eggs *often enough*. And often enough, for sperm, is not very often. But at the other end of the spectrum some devices fulfil their functions all of the time when functioning normally in normal conditions. The heart provides a good illustration; it pumps blood all the time. Less reliable hearts were selected out. Some items thus have *reliability* vis-à-vis a certain effect as their function. The heart clearly does; the heart is for *reliably* pumping blood.

Functional items often have a plurality of functions. One item may be intended, designed, or selected to do many things. It may have more than one causal-role capacity. Most do. It is thus misleading to speak of 'the' function of an item. When it comes to etiological functions and biological items in particular,

¹¹ Since an item may fail to fulfil its function even while in normal conditions while functioning normally, it is a mistake (too easily made) to suppose that failure to fulfil a function implies that the item is either not in normal conditions or not functioning normally (malfunctioning).

there is often a plurality at ‘horizontal’ levels on the ‘vertical’ dimension. An example will show what I mean. The tongue is for eating; it moves food around as we chew, and also assists in swallowing. But the tongue is also for talking; it helps form word sounds in conjunction with other parts of our mouth and vocal tract. These are all functions of the tongue at the same ‘horizontal’ level. By doing all of these things, the tongue assists in fulfilling other functions, such as getting enough food and nutrients into the bloodstream so as to nourish other parts of the body, so as to remain healthy and fit, so as to assist in dealing with the creature’s natural habitat, so as to find more food, avoid predators, find mates, and so on, so as to live long enough to survive to reproduce. All of these other more distal effects of the tongue are among its functions. On this ‘vertical’ dimension, the tongue has many functions. (And so my title is a bit misleading. I argue for *a* function of comprehension and not for its *one and only* function.)

Items also lose e-functions over time (the appendix), acquire additional e-functions over time (feathers), and even lose one e-function only to acquire another (limbs). When we speak of the e-functions of an item, we often mean current functions (Millikan 1984, 2002; Godfrey-Smith 1994; Schwartz 2002). Origins are one thing; functions (sometimes) are another.

To sum up, there are (at least) four kinds of functions: designed, intention-determined, causal-role, and etiological functions. They all differ notionally and extensionally, though in particular cases two or more may overlap. When it comes to items with etiological functions, *ceteris paribus* they fulfil their functions in normal conditions when functioning normally. Normal functioning is thus operating or working in a manner that, in normal conditions, leads to function fulfilment. If you know an etiological function of an item, you can thereby partially individuate and explain normal functioning in terms of that function. And some items—like the heart—have reliability vis-à-vis an effect as a function. Normal functioning, for the heart, is thus individuated and understood, in part, in terms of reliably pumping blood.

3. EPISTEMIC ENTITLEMENT

I now use the etiological notions of function and normal functioning to explicate (a kind of) epistemic entitlement (cf. note 3). I argue that entitlement attaches to beliefs in virtue of the normal functioning of the belief-forming process when the process has forming true beliefs reliably as an etiological function.

First, positive epistemic standings are goods, successes, fulfilments, or achievements understood in terms of promoting true belief and avoiding error. Warrants are positive epistemic standings that epistemically support or ground beliefs. Warrants are positive epistemic standings, and so are themselves goods, successes, fulfilments, or achievements understood, at least in part, in terms

of promoting true belief and avoiding error. Any explication of a kind of warrant must explain why the warrant is a success, good, fulfilment, or achievement understood in terms of promoting true belief and avoiding error.

Second, function fulfilment is a success or good. The function sets a standard or norm. Fulfilling the function achieves that standard. Meeting the standard is success or fulfilment, and as such is a good for the item with the function. And failure to fulfil the norm is just that, failure. Normal functioning is also a fulfilment. Functioning normally is functioning (operating, working) the way the item is supposed to operate. Think of the heart. Beating normally is the way the heart is supposed to beat. In beating that way, it pumps blood. Beating normally fulfils a standard for the heart. And failure to operate normally is just that, a failure. Normal functioning is a success or fulfilment for the functional item, and as such is a good for the item. Using 'achievement' broadly, function fulfilment and normal functioning are both achievements for functional items.

Despite their explanatory interrelation, function fulfilment and normal functioning are *distinct* achievements. Recall that you can fulfil your function without functioning normally, and vice versa. A heart in a sterile dish may function (work, operate) normally without blood passing through. A car up the lift may function just fine without taking you where you want to go.

Third, for any belief-forming system or process whose etiological function is to form true beliefs reliably, *ceteris paribus* it will form true beliefs reliably in normal conditions when functioning normally. This follows automatically from the logic of etiological functions. Normal functioning for such a process is thus functioning in a way that, in normal conditions, reliably produces true beliefs. Normal functioning for any such system is then individuated and understood in terms of reliability and truth. Again, all of this follows from the logic of e-functions, as just explained.

Putting two and two together, we arrive at an account of epistemic entitlement. Entitlement is a kind of warrant that does not depend on reasons or evidence available to the subject. Entitlements are psychologically and developmentally more primitive kinds of warrants. *Warrants* generally are *fulfilments, successes, or goods* understood, at least in part, in terms of promoting *true* belief and avoiding error that ground or support belief. *Normal functioning* is an *achievement, success, or good* understood in terms of *function fulfilment*. So for any belief-forming process whose e-function is *reliably producing true beliefs*, normal functioning for that process is an achievement understood, at least in part, in terms of promoting true belief and avoiding error. When the belief-forming process functions normally, it epistemically grounds or supports the belief it thereby causes or sustains. A kind of warrant thus attaches to belief when formed on the basis of a normally functioning belief-forming process whose function is to form true beliefs reliably. Since this warrant does not depend on reasons or evidence

available to the subject, it is (a kind of) entitlement. A normally functioning belief-forming process that has forming true beliefs reliably as a function confers entitlement on the beliefs it normally causes and sustains. I have established my first thesis, an account of (a kind of) entitlement.

What does this account imply about the massively deceived brain-in-a-vat? In such a case, the standard belief-forming processes are all functioning normally, but they are not reliable. Envatted, the subject's belief-forming process may all function (operate, work) normally, even while not in normal conditions. They thus enjoy (a kind of) entitlement just as if the subject were not envatted. The brain is like a car up the lift for inspection. The subject's belief-forming processes are operating just fine, even though they are not producing and sustaining true beliefs at, so to speak, sixty-five miles an hour.

And what does this account imply about the beneficiary of an accidentally acquired, merely reliable belief-forming process, a process the subject knows absolutely nothing about (that he has it, that it is reliable, and so on)? Merely accidentally acquired, the process has no function. It's *just* a mutation. Nothing then counts as normal functioning for the process. Though the subject's beliefs so caused enjoy the epistemic good of being mostly true, they lack the achievement conferred by function fulfilment, for it lacks a function, and they lack the achievement conferred from normal functioning, for nothing counts as normal functioning.

Space does not allow me the luxury of responding to familiar objections to etiological accounts of epistemological properties. I hope I have said enough here to make my account of entitlement *prima facie* plausible. Elsewhere I answer some of the more prominent objections.¹²

Perceptual beliefs paradigmatically enjoy entitlement, for (most) perceptual belief-forming systems have reliably forming true beliefs as an etiological function. I argue that comprehension-with-filtering-based beliefs enjoy entitlement for the very same reason; comprehension with filtering has forming true beliefs reliably as an e-function. When our capacity to comprehend and filter functions normally, it meets a standard understood in terms of promoting true belief and avoiding error. That is why our comprehension-with-filtering based beliefs enjoy *prima facie pro tanto* entitlement. (As I noted, this a good place to take a break.)

4. AN ETIOLOGICAL FUNCTION OF ASSERTION

But before arguing that *comprehension-with-filtering* has *reliably* causing and sustaining true beliefs as an e-function, I first argue that *assertion* has inducing (and sustaining) *true* hearer belief as an etiological function. (Recall I am using 'assertion' broadly.) I take up assertion first for two reasons. Making the

¹² See, for example, Graham (unpublished).

case for assertion practically makes the case for comprehension-with-filtering; if the argument is plausible for assertion, it will be even more plausible for comprehension-with-filtering. And most of the relevant discussion in the literature so far has focused on assertion instead of comprehension.¹³ As a result, the question whether *comprehension-with-filtering* is for *reliably* forming true beliefs is apt to be confused with whether *assertion* has that function. Discussing assertion first should head off potential misunderstanding.

To show that assertion has inducing true hearer beliefs as one of its etiological functions, I begin with an argument from C. A. J. Coady intended to demonstrate the reliability of any practice of reporting. Though Coady's argument unsurprisingly falls short, it reveals an important fact about assertion.

Coady conducted a thought experiment involving a community of Martians. 'Let us suppose for the moment,' he writes,

that they have a language we can translate (there are difficulties in this supposition . . .) with names for distinguishable things in their environment and suitable predicative equipment. We find however, to our astonishment, that whenever they construct sentences addressed to each other in the absence (from their vicinity) of the things designated by the names, but when they are, as we should think, in a position to report, they seem to say what we (more synoptically placed) can observe to be false. (1992: 85)

Coady then asked whether any hearer would have any reason to rely on what a speaker says in such a situation. He replies that

any Martian has four powerful reasons for not relying on what others appear to be telling him: (i) he finds their 'reports' false whenever he checks personally on them, (ii) he finds reliance upon them consistently leads him astray in practice, (iii) he finds himself utterly unreliable in what he tells others and it is, at least in part, possible that he is not atypical, (iv) others often give chaotically different reports on those matters beyond his checking. (1992: 87)

Coady inferred there would then be no practice of making and accepting reports. Given what each Martian can expect of the other,

it is . . . very hard to imagine the activity of reporting in anything like its usual setting with the Martians, for there would surely be no reliance upon the 'reportive' utterances of others. [T]he Martian community cannot reasonably be held to have the practice of reporting . . . (1992: 87)

Coady concluded that any practice of reporting must be reliable.

I think Coady missed a premise. With the premise he needs, Coady's argument goes like this:

- (1) If everyone in a community, when making reports, always makes false reports, then no one will have any reason for relying upon and accepting those reports. They have 'four powerful reasons not to do so'.

¹³ E.g. McDowell (1980), Millikan (1984), Sperber (2001), Origgi and Sperber (2000), Faulkner (2000), and Schmitt (2006).

- (2) A central reason (motive) for making reports is to get others to believe them.
- (3) But if no one accepts those reports, then speakers will have no reason to think that making those reports will lead hearers to believe them. (This is the missing premise.)
- (4) So if no one accepts reports, speakers will stop making them.

In short: if all falsehoods uttered, then no acceptances; if no acceptances, then no reason to report; if no reason to report, no reporting. Contraposed: if reports, then not all reports are false. Hence: if reports, then truths.

In 'The Reliability of Testimony' (Graham 2000c) I showed that this argument falls short of establishing the *reliability* of *any* practice of reporting. At best it shows that reporting practices exist because not all of the reports made are false. And 'not all are false' surely does not entail 'most are true'.

Though this argument does not show what Coady hoped, it does bring out a very important fact about linguistic practices: speakers and hearers both need some reason (motive) to participate. Speakers, presumably, benefit in some way by affecting hearers. If hearers receive no benefit from being so affected, they will probably stop responding in the desired way. So unless hearers get something out of accepting reports, they will not accept them. And if they will not accept them, speakers will not benefit by making them. Then they will not get made. Hearer benefits (partly) explain speaker production.

Ruth Millikan exploits this very fact when arguing that various linguistic devices have etiological functions. She has argued in general that (i) language devices produce effects that interest speakers often enough to encourage continued replication by those speakers *only if* hearers produce hoped-for responses often enough, and (ii) that hearers will continue to produce those hoped-for responses often enough *only if* the results are of interest to hearers (Millikan 2004: 25). Language devices exist, as a matter of fact, only if their use benefits both speakers and hearers.

Applied to (declarative sentences in) the indicative mood, the device we paradigmatically use to make assertions, she has argued that one of its central functions is to produce true beliefs in the propositional content(s) of the utterance, for it is the effect that, often enough, benefits hearers. Truth about relevant subject matters clearly benefits hearers. Dan Sperber says that:

From the point of view of receivers, communication, and testimony in particular, is beneficial only to the extent that it is a source of genuine (and of course relevant) information. Just as in the case of individual cognition, there may be cases where biases in communicated information are beneficial (think of exaggerated encouragement or warnings, for instance), but these cases are marginal. (2001: 404)

Though inducing false beliefs sometimes benefits speakers, and it may also sometimes benefit hearers, if hearers never got true information from assertive utterances of declaratives in the indicative mood, they would quit forming beliefs

in response. And then speakers could not even get across false information when they wanted to. (This is what Coady's thought experiment brings to the fore.) And if hearers quit forming beliefs in response, then speakers would quit making those utterances. There would be no point. True hearer belief thus explains why the indicative mood persists. True hearer belief is thus a function of the indicative mood:

If no true beliefs ever resulted from hearer interpretations of indicative sentences, it is clear that indicative syntactic patterns would cease to be used first by hearers and, as a result, by speakers in the ways they now are. . . . The focused stabilizing function of the indicative mood is thus the production of a true hearer belief. 'The function' of the indicative mood is to convey information. (Millikan 1984: 58–9)¹⁴

Millikan calls this function a stabilizing one, for it 'encourages speakers to keep using the device and hearers to keep responding to it with the same (with a stable) response' (2005: 94).¹⁵

Since the default, standard use of a declarative sentence in the indicative mood is to assert, and since the same reasoning Millikan offers applies to assertive communication generally, I conclude that an *etiological* function of assertion is to induce true belief in hearers. Truth is a price speakers pay that keeps assertion in play. Though Coady failed to establish the necessary reliability of any practice of reporting, he unearthed an important fact that helps show that inducing true belief in hearers is an etiological function of assertion.

5. ALTERNATIVE FUNCTIONS

Some will resist; assertions may regularly *cause* true beliefs—assertion may even as a matter of fact cause *mostly* true beliefs—nevertheless inducing true belief is *not* what assertion is *for*. In this section I reply to four counter-arguments. Responding to these should help persuade the unpersuaded. The last counter-argument will also take us from the function of assertion to the function of comprehension-with-filtering.

¹⁴ This argument appears throughout Millikan's writings (cf. 1984: 30–2, 52–60, 2004: 25–6, 2005: 14, 16, 44–5, 57–8, 90, 94–5, 153, 157, 169, 172).

¹⁵ Though Dan Sperber (rightly) doubts assertive communication will be as reliable as perception, he agrees with Millikan on evolutionary grounds: 'For communication to stabilize within a species, as it has among humans, both the production and the reception of messages should be advantageous. If communication were on the whole beneficial to producers of messages (by contributing to their fitness) at the expense of receivers, or beneficial to receivers at the expense of producers, one of the two behaviors would be likely to have been selected out, and the other behavior would have collapsed by the same token (incidentally, there are exceptions, in particular in inter-species communication). In other words, for communication to evolve, it must be a positive-sum game where, in the long run at least, both communicators and receivers stand to gain' (Sperber 2001: 403). For further agreement and discussion of the stability of truth-telling, see Sober (1994a).

The first counter-argument focuses on what the speaker has in mind when making her assertion. Call this the *argument from occasion specific intentions*. The argument goes as follows. First, a function is a goal or purpose. Second, the function of a particular assertion on a particular occasion is the goal, purpose, or intention of the speaker on that occasion. Third, speakers have all sorts of purposes on various occasions. Sometimes they intend that their interlocutors form true beliefs, sometimes they only intend shared belief, sometimes they don't intend belief at all, sometimes they intend false belief, and sometimes the truth status is just irrelevant. True hearer belief is thus not the intended effect on each and every occasion. Hence true hearer belief cannot be what assertion, as a type, is for. Assertion is for whatever speakers use it for.

Though I think this argument may be in the mind of many readers, it is rather easy to rebut. From the fact that a particular speaker does not fail at fulfilling *his* aim when making a false assertion, it does not follow that the *etiological* function of the linguistic device he used performed *its* function.

Consider the heart. Though its function is to pump blood, a malicious surgeon may use one as a paperweight. Or consider a bird's wing. Its e-function is to enable flight, though you can cut it off and use it as a fan. Cut off from the bird, it clearly does not perform its function, even if it keeps you cool. A speaker 'can use a language device, just as he can use a tool or one of his body parts, to perform a function or to serve a purpose that is not, however, the language device's *own* function' (Millikan 1984: 52). E-functions and i-functions are simply different things. So from the fact that a speaker does not intend to convey truth on a particular occasion when asserting, nothing follows about the e-functions of assertion.

There is, however, a connection between speaker purposes and the function of assertion. Speaker purposes must be satisfied often enough, or they will stop making assertions. So, in a clear sense, the fulfilment of speaker intentions in general is an etiological function of assertion. But the Coady–Millikan argument places a constraint on speaker purposes. Speakers must purpose true belief in hearers *often enough* for true belief to be a function of assertion. Here again is Millikan:

But although the stabilizing function of a language device is independent of the purpose of the particular speaker who utters it, it is not independent of speaker purposes in general. The survival of a public language element without change of function must depend upon their being a critical mass of occasions upon which speakers and hearers use the element such that it performs its stabilizing function. (Millikan 1984: 53)

And so long as speakers purpose true belief *often enough*, hearers will respond to assertions with belief *often enough* for speakers to keep making assertions. So speakers satisfy their purposes in general when making assertions because they purpose true hearer belief often enough. Speaker benefits are thus intertwined with hearer benefits, the benefit of true belief in particular.

The next counter-argument exploits the standard Gricean view of communication. Call it the *argument from successful communication*. On the Gricean view, communication succeeds or occurs when the hearer recognizes the speaker's communicative intention. A communicative intention is the speaker's intention that the hearer recognize that the speaker intends that the hearer form a belief, answer a question, follow a command, and so on. When the hearer recognizes the communicative intention communication occurs, even if the hearer does not form the belief, answer the question, or follow the command.

Here is Stephen Levinson from his classic textbook:

[C]ommunication involves the notions of intention and agency, and only those inferences that are openly intended to be conveyed can properly said to be communicated. . . . Communication consists of the 'sender' intending to cause the 'receiver' to think or do something, just by getting the 'receiver' to recognize that the 'sender' is trying to cause that thought or action. So communication is a complex kind of intention that is achieved or satisfied just by being recognized. [. . .] Attaining this state of mutual knowledge is to have successfully communicated. (Levinson 1983: 15–16)

Kent Bach puts the point more succinctly:

Communicative success is achieved if the speaker chooses his words in such a way that the hearer will, under the circumstances of utterance, recognize his communicative intention. (. . .) An act of communication is successful if whoever it is directed to recognizes the intention with which it is performed. In short, its fulfilment consists in its recognition.

(Bach 2006: 153)

In a simple case, in assertively uttering *U*, the speaker intends that the hearer recognize the speaker's intention that the hearer believes the propositional content(s) of *U*, and perhaps various implicatures. Communication 'succeeds' if the hearer recognizes the intention. A speaker can thus communicate when asserting *P* whether or not *P* is true, and whether or not his interlocutor comes to believe *P*. That is simply how communication works.¹⁶

With this view of 'communicative success' in hand, we can state two related arguments for the conclusion that the function of assertion isn't inducing true hearer belief. The first focuses on the transpersonal communicative process. The second focuses on the speaker's side of the equation. I'll state and discuss the first before turning to the second.

So first think of communication as a transpersonal process: if the hearer correctly recognizes the speaker's communicative intention, then communication has occurred. If uptake did not occur, neither did communication.

¹⁶ A casual view of hearer recognition of speaker's communicative intentions might suppose that it consists in the (true) conscious explicit belief that the speaker intends the hearer to recognize his (the speaker's) informative intention. I seriously doubt whether the recognition consists in such a belief. One reason for this is that young children comprehend assertions—and thereby recognition communicative intentions (if the Gricean account is correct about communication)—but they do not have the capacity to form such beliefs. See also Hunter (1998).

Here is the counter-argument. First, function fulfilment is a success; failure to fulfil a function is just that, failure. So we can derive the function of communication from what counts as communicative success and failure. Third, communicative success involves recognition of the speaker's communicative intention, and does not require forming belief in the content of the assertion, let alone forming true belief. So if communication can succeed (occur) without conveying true belief, true belief cannot be what it is for; inducing true hearer belief cannot be a function of assertion.

This argument is a non-starter. For it conflates the *ontological* question of what communication *is* (what it consists in, what must happen for it to occur) with the *teleological* question of what it is *for*. The relevant sense of 'success' and 'failure' in the third premise is ontological or constitutive. But the relevant sense of 'success' and 'failure' in the first premise is teleological or purposive, where functions may go unfulfilled. To say what assertive communication is *for* requires pointing to some effect of successful communication. This argument is silent on that issue. Nothing in the argument shows that inducing true beliefs is not what communication is for.

This problem, however, may be avoided if we formulate the argument in the second way mentioned, in terms of the speaker's communicative act itself and its function. Here we ask what the function of the speaker's communicative act happens to be, an act that may fail in fulfilling its aim. The counter-argument would then go as follows. First, function fulfilment is a success; failure to fulfil a function is just that, failure. So we can derive the function of an assertive communicative act from what counts as success and what counts as failure. Third, the speaker's assertive communicative act—his assertion—succeeds when the hearer recognizes the speaker's communicative intention. The etiological function of assertion would then be to induce hearer recognition of the speaker's communicative intention. And recognizing this intention does not require that the hearer believe what the speaker asserted, let alone require that the hearer believe something true. So inducing true hearer belief cannot be what assertion is for. Paraphrasing Paul Faulkner, *when a hearer fails to form a true belief from communication, there is no sense in which the speaker has failed to communicate* (2000: 587).

Though this version of the argument from communicative success focuses on assertion and its effects, and thereby avoids the rejoinder just made to the first version of the argument, it fares no better. Simply put, it doesn't focus on effects that explain why assertion as a type persists. An etiological function is an effect of a type that occurs often enough that explains why the type persists. This argument does not address this issue. All it points out is that a speaker, when making an assertion, intends that the hearer recognize his communicative intention, and when the hearer does, the speaker has communicated. But what is all of this for? What effect explains why speakers make assertions? What effect explains why speakers publicize communicative intentions? What effect explains

why hearers even bother to recognize communicative intentions? The argument, though arguably correct about the mechanics of assertive communication, does not address the etiological question of what assertion is for.

Levinson himself is clearly aware of the fact that an *account* of how communication works is one thing, and what linguistic devices are *for* is another. Well after arguing for a Gricean account of communication, he notes that

[since] nearly all of the world's languages have the three basic sentence types: imperative, interrogative and declarative . . . one might argue . . . on the grounds that these seem to be used paradigmatically for ordering, questioning, and asserting . . . that they recur in the languages of the world because humans are, perhaps, specifically concerned with three functions of language in particular—the organization of other persons' actions, the eliciting of information, and the conveying of information. (1983: 40)

The mechanism of communication is one thing; its purpose is another. Griceans may be right about how communication works, while Darwinians (so to speak) are right about what communication is for.¹⁷

Another counter-argument comes from a paper by Gloria Origgi and Dan Sperber entitled 'Evolution, Communication, and the Proper Function of Language' (2000). They target Millikan's view that the indicative mood is for inducing true hearer belief. They argue that its function is to provide evidence to assist the hearer in recognizing the speaker's communicative intention. The 'function of linguistic utterances', and so the function of assertions in particular, is not to affect hearer beliefs in the content of what is asserted, but rather to provide 'highly precise and informative evidence of the communicator's intention' (Origgi and Sperber 2000: 163).

The following passage contains Origgi and Sperber's reasoning. I have added emphasis to draw attention to a key premise:

Our disagreement with Millikan has to do with the level of processing at which linguistic devices elicit *the reliable response to be identified as their direct proper function*. For Millikan, this *reliable response* is to be found at the level of belief or desire formation, or even at the behavioral level in the case of compliance . . . What is the alternative? Linguistic

¹⁷ The *argument from occasion specific intentions* and the *argument from communicative success* are clearly related. The first says the function of a particular assertion is whatever intention the speaker has in mind in making the assertion. The second says the function of communication is the recognition of a particular intention on the part of the speaker, an intention that is always present when an assertion is made, always recognized when communication occurs. The function of assertion on both arguments is just the fulfilment of a speaker's intention. As Millikan has noted (2005: 28), tying function to intention is a dominant, traditional philosophical view about how communicative functions are determined. This view has wide appeal outside of philosophy as well. Witness Michael Tomasello: 'The notions of communicative intention and function are correlative. Someone uses a piece of language with a certain communicative intention, and so we may say that that piece of language has a certain function' (2003: 3). But since we know that functions from histories are one thing and functions from intentions are another, we know that neither of these arguments touches the thesis that the etiological function of assertion is to induce true hearer belief. For functions from intentions, look to intentions. For functions from histories, look to histories.

devices produce *highly reliable responses*, not at the level of the cognitive outputs of comprehension such as belief or desire formation, and even less at the level of behavioral outputs such as compliance, but at an intermediate level in the process of comprehension. Linguistic comprehension involves, at an intermediate and largely unconscious level, the decoding of linguistic stimuli that are then used as evidence by the hearer, together with the context, to arrive inferentially at the speaker's meaning . . . Linguistic devices have proliferated and stabilized because they cause these *highly reliable cognitive responses* at this intermediate level. Linguistic devices provide speakers and hearers with informationally rich, highly structured, and reliably decoded evidence of speaker's meaning.

(Origg and Sperber 2000: 163)

Call this the *argument from highly reliable effects*. The argument goes like this. First, etiological functions are highly reliable effects. Second, the highly reliable effect of assertions (and other linguistic devices and speech acts) is not true hearer belief but the providing to hearers of evidence for the speaker's communicative intentions. Hence providing evidence of the speaker's communicative intention, or perhaps the hearer's decoding and recognizing that evidence (Origg and Sperber speculate on where to locate 'the function' in the causal chain) is the etiological function of linguistic devices.

The argument makes two mistakes. First, as we have seen, etiological functions are not necessarily highly reliable effects. True, pumping blood is a highly reliable effect of the heart, and the function of the heart is to pump blood. But this is not true of all items with etiological functions. Recall sperm. Millikan would be the last person to identify functions with highly reliable effects. Though some items do have reliability as an aspect of their function (like the heart), some clearly do not. You thus cannot infer function from highly reliable effect. Indeed, some highly reliable effects are not even functions. The heart makes noise just as frequently as it pumps blood, and making noise is not one of its functions.

Second, even if providing evidence of the speaker's communicative intention is *an* etiological function of a linguistic device, it does not follow that it is the *only* etiological function. Recall that items with etiological functions often have a number of functions, both horizontally and vertically. Just as the heart assists in removing wastes and circulating nutrients by pumping blood, and it pumps blood by beating at a certain rate in a certain way, so too assertion may induce true hearer belief by providing evidence of speaker's intentions. Assertion, like any other functional item, can have more than one function. Linguistic devices may have proliferated and stabilized (in part) because they provide evidence of speaker's intentions, but they have also proliferated and stabilized by having effects on hearers that explain why hearers care to recover speaker's intentions from the evidence provided in the first place. Origg and Sperber have provided no reason for thinking that inducing true hearer belief cannot be *an* etiological function of assertion. At best they are correct that inducing true hearer belief is not its *only* function. And so this argument too fails to undermine the thesis that inducing true hearer belief is an etiological function of assertion.

Though I have already expanded at some length defending this thesis, I would be amiss to not address one last counter-argument. The argument derives from discussions of the evolutionary origin of the human language faculty. It is widely held in those discussions that language arose in order to facilitate social bonding and social interaction, so that our primate evolutionary ancestors might operate effectively in larger groups, where larger groups promote higher levels of fitness. Robin Dunbar argued for this view at length in his book *Grooming, Gossip, and the Evolution of Language* (1996). Here is Robbins Burling stating the older view of the origins of language in his book *The Talking Ape*:

There was a time, not so many decades ago, when anyone who thought at all about the evolution of language took the reasons for its selection to be more or less self-evident. We use language today for every sort of practical purpose, to exchange information, to coordinate our daily lives, and to instruct the young, so it only seemed reasonable to suppose that, from its earliest days, hunters and gathers used language for much the same purposes. (Burling 2005: 181)

He goes on to embrace the ‘social interaction’ view:

I share a growing consensus that our intelligence evolved primarily as a means for dealing with other individuals. Ours is primarily a social intelligence, and language is one product, and one part of, that intelligence. In preurban societies, and even our own, language in its most delicately nuanced form is used, not so much for basic subsistence tasks, as for establishing, maintaining, and refining social relationships. It is when dealing with people, not with material objects, that we call on our richest linguistic resources.

(Burling 2005: 184)

And so we once thought that language arose for ‘exchanging information’ but we now think language arose for ‘dealing with other individuals’. Communication, and assertion in particular, is thus not for inducing true hearer belief, but rather for interacting with others and maintaining social bonds. So instead of focusing on more proximal effects of assertion (such as providing evidence of speaker intentions or the actual recognition of such intentions), this argument focuses on the more distal effects of inducing belief through communication. Assertion is thus for social bonding, not for conveying true information. Call this the *argument from social interaction*.

The error in this argument should be easy to spot. There is no opposition between social bonding and conveying information as etiological functions for assertion. A closer reading of Dunbar and Burling (see also Pinker and Bloom 1990; Dessalles 2007; Hurford 2007) reveals no such opposition at all. In fact, Dunbar holds that language conduced bonding and so the forming of larger groups *by* conveying information about social matters. Social bonding is simply higher up in the vertical hierarchy of language functions. Conveying true information is one etiological function of communication:

language evolved to facilitate the bonding of social groups, and . . . it mainly achieves this aim by . . . the exchange of socially relevant information. (Dunbar 1996: 120)

language evolved to exchange information about other individuals in the group . . . language evolved to exchange social information. (Dunbar 1996: 142)

Language . . . is very much a social tool . . . it allows us to exchange information relevant to our ability to survive in a complex, constantly changing social world . . . (Dunbar 1996: 170)

And so the old view is that language evolved primarily to convey information about where and how to find food and avoid danger. The new view is that it evolved in addition, if not primarily, to convey information about social matters to build groups where building groups, in turn, has obvious implications for survival. Burling concludes:

Thanks to the eagerness of people to display their knowledge, we are able to learn important things from conversation and it is people that we learn most about. Since it is maneuvering through the social system that gives us the most difficult challenges of our lives, the information we gain by learning how others have behaved is extremely valuable. It is fortunate for the hearer, then, that we are at least as eager to pass on information as to receive it. (Burling 2005: 196)

Either way, language originated, in part, to convey true information.

Summing up, assertive utterances provide evidence of speaker's intentions. Hearers recognize speaker's intentions and thereby comprehend speaker's assertions. Comprehension in turn induces true belief, often enough, and thereby provides an obvious benefit to hearers. Conveying true belief builds social bonds, where building bonds conveys benefits to both speakers and hearers. Conveying true information is one effect of assertion that clearly explains why assertion, as a type, persists. 'The immediate effect of language when used in standard conditions, in other words, its proximal function, is . . . to convey . . . the information and knowledge possessed . . . by the speaker' (Dessalles 2007: 315).¹⁸ (In need of a break? This would be another good place to pause.)

6. COMPREHENSION, RELIABILITY, AND FILTERING

Absent counter-considerations, we take comprehension at face value. My topic is the epistemology of this process. What kinds of warrant do beliefs

¹⁸ John McDowell has argued that conveying knowledge is the function of communication (1980). Though obviously congenial to my point of view, McDowell's paper lacks (a) an explication of functions, especially etiological functions and how they differ from functions from intentions, (b) a discussion of normal functioning, (c) the distinction between a one-off effect as function, and reliability vis-à-vis that effect as function, (d) an extended discussion of why conveying knowledge should be the function, and (e) an explanation of the relevance of the function of communication to the epistemology of beliefs based on comprehension. McDowell also focuses on knowledge from comprehension, and does not isolate entitlement as a distinctive epistemic achievement.

based on comprehension-with-filtering enjoy? I argued a priori in section 3 that beliefs based on a normally functioning process that has forming true beliefs reliably as an etiological function enjoy *prima facie pro tanto* entitlement. Beliefs based on comprehension-with-filtering will thus enjoy entitlement provided comprehension-with-filtering has forming true beliefs reliably as an etiological function. Does comprehension-with-filtering have that function?

I have just argued that the usual causal precursor to comprehension—assertion—has inducing true hearer belief as an etiological function. If assertion is for producing true belief, and it does so by triggering comprehension states, so comprehension states are also for inducing true hearer belief. If not enough of the beliefs formed on the basis of comprehension were true, then hearers would not be moved to comprehend assertions or transition from comprehension to belief in the first place. The argument of section 4 and its defence in section 5 carries over automatically to comprehension. And since the transition from comprehension to belief is not intentional in the way assertions are, the first three counter-arguments just discussed do not touch the extension of the argument from assertion to comprehension. Comprehension has inducing *true* beliefs as a function.

But what about inducing true beliefs *reliably*? Does comprehension have producing true beliefs *reliably* as a function?

I am happy to grant the possibility that comprehension taken alone, comprehension without filtering, may fail to have inducing true beliefs reliably as a function. I am happy to grant the possibility for my thesis does not require it. My thesis is that comprehension-*with-filtering* has forming true beliefs reliably as a function, *not* comprehension *neat*, comprehension *taken alone*. It's the filtering, or so I argue, that is for producing a sufficiently high truth ratio. On the other hand, if it turns out that comprehension neat reliably induces true beliefs, and that is why we do it, then comprehension as such would confer entitlement when functioning normally. My argument here would then establish a higher degree of entitlement for beliefs based on comprehension-with-filtering.

But before I argue that comprehension-with-filtering is for reliably inducing true belief, let me make one point very clear. I am not arguing a priori that comprehension-with-filtering is necessarily reliable. My argument is broadly empirical. What's a priori is my account of entitlement. Its application to comprehension-with-filtering is empirical (cf. Goldman 1986).

So why should we think comprehension-with-filtering induces true beliefs *reliably*, and why should we think it has inducing true beliefs reliably as an etiological *function*? The answer turns on understanding why we filter, what filtering is for.

Recall the logic of the argument derived from Coady and Millikan. Speakers assert for some benefit. Hearers in turn form beliefs for some benefit. Unless hearers formed true beliefs often enough, they would stop forming beliefs as a response to speaker assertions. Speakers would no longer benefit by making assertions. Assertions would then no longer get made. Inducing true hearer belief is thus a stabilizing function of assertion. So far I have emphasized how hearers benefit: they form true beliefs. I have not discussed how speakers benefit, benefits that explain why we are so disposed to make assertions in the first place.

There are many ways speakers might benefit by making assertions. By providing useful and relevant information, speakers gain status and build coalitions (Dessalles 2007), build social bonds (Dunbar 1996; Burling 2005), enhance cooperative activities (Tomasello 2008), and even attract reproductive partners (Burling 2005). And the argument so far reveals that speakers acquire these benefits by providing true information, at least often enough. Some of these benefits explain the origins of informative speech and communication. Some explain its current shape. Providing true information leads to speaker benefits that historically and currently explain why assertions get made, and so explains why assertion as a type persists.

So we benefit both as speakers and hearers in a variety of ways from communication, where the passing along of true information plays a vital role in accounting for all of this. But despite the utility of true information, how often do we actually present the truth? How honest are we? Don't we often deceive? What about the actual frequency of cheats and liars? Don't speakers often appear to cooperatively provide useful information that turns out to be bogus? And can't speakers sometimes reap the benefits associated with making assertions but without paying the price of providing true information? Perhaps speaker benefits on occasions may be best achieved not by providing true information, but by distorting the truth, or even saying what is clearly false (Sperber 2001). And won't the cost to speakers of acquiring true and useful information, and the cost of sharing such information, sometimes serve as a motive for trying to reap the benefits of being a good information provider by providing apparently genuine information that is not the real thing (Dessalles 2007)? Won't conflicting motives and the cost of acquiring information lead speakers, on occasions, to say what is not the case? Isn't some degree of misinformation guaranteed by the logic of the argument? Dan Sperber concludes that

a significant proportion of socially acquired beliefs are likely to be false beliefs, and this not just a result of the malfunctioning, but also of the proper [normal] functioning of social communication. . . . [T]he cognitive manipulation of others is one of the effects that makes the practices of testimony and argumentation adaptive. This contributes to explaining why these practices have evolved and stabilized among humans.

(Sperber 2001: 402)

Some degree of misinformation is thus bound to result from the very logic of the argument that shows that assertion has providing true information as (one of) its etiological function(s).

On the other hand, hearers do not sit idly by as speakers mislead or otherwise misinform. Hearers develop countermeasures. Hearers develop capacities for filtering out reliance on fake information providers. Indeed, for communication to stabilize, hearers must develop such filters. Here again is Sperber:

if communication has stabilized among humans, it must be that there are ways to calibrate one's confidence in communicated information so as that the expected benefits are greater than the expected costs. (Sperber 2001: 406)

One important filter consists in detecting cognitive conflicts: coherence checking. Sperber suggests that

coherence checking—which involves metarepresentational attention to logical and evidential relationships between representations—evolved as a means of reaping the benefits of communication while limiting its costs. (Sperber 2001: 410)

Jean-Louis Dessalles agrees with the challenge and its solution:

Let us assume that whatever benefit speakers derive will increase with the salience of the situations they speak of. This makes for a strong temptation to exaggerate or even to tell lies. The risk run by a hearer is that of affording an undue benefit to the speaker. Against this risk, hearers have two strategies, trivialization through comparison with already known situations and detection of inconsistencies. Both of these rely on faculties with which natural selection has endowed us as human beings, among which are a sense of probabilities and the ability to detect cognitive conflicts. Each of these conversational mechanisms has thus much the same function as the other, that is avoidance of cheating in communication. (Dessalles 2007: 331)

hearers . . . run the risk of being cheated in verbal interactions . . . The original role of the detection of cognitive conflicts can be made sense of as the attempts of hearers to guard against lying. (Dessalles 2007: 330)

Though there are incentives for speakers to cheat or mislead, we develop filters for screening out their influence. Coherence checking is one. We encountered other filters in the first section. We deploy all of these and more.¹⁹

¹⁹ Speaking of cooperative communication in general, and not just human linguistic communication, Marc Hauser writes: 'The last 20 years or so in biology has taught us . . . that whenever individuals assume different roles in an apparently cooperative venture, there are opportunities for individuals to cheat, reaping the benefits without paying the costs' (2000: 141). It's a general strategy. He continues: 'some animals have the capacity to detect mismatches between a signal's common function and the way it is currently being used. These animals can therefore detect cheaters, liars, and other self-interested deceivers. They have the mental tools to see through deceit' (2000: 156). We are animals like these.

Filtering thus plays a very important role; it ensures the reliability of beliefs formed on the basis of comprehension, for misleading or untrustworthy assertions often get filtered out. We're not so easy to mislead:

hearers have effective ways of assessing the quality of information supplied to them and . . . this makes it difficult to mislead them. (Dessalles 2007: 332)

the only signals that natural selection can favour are the reliable ones. With language . . . the reliability of the message is . . . guaranteed by . . . its resistance to hearer's assessments. . . . It is not . . . easy to tell lies, even with words, as hearers test the logical consistency of what they are told. (Dessalles 2007: 331)

Furthermore, filtering not only dampens the possibility of accepting a false report, it also provides an incentive for speakers not to cheat in the first place. For cheaters often get caught. Indeed, a good deal of the social gossip hypothesis is driven by the idea that one very useful kind of information we receive from communication is information about which speakers tend to cheat (Dunbar 1996). Given the cost of getting caught, nature has even implanted in speakers a disposition to stick to the facts:

[We have] an emotional warning whenever we are tempted to lie. Our own nervousness cautions us to be careful, and the nervousness may be visible enough to make others suspicious. We are not as good at lying as evolutionary theory and sheer self-interest might lead us to expect, but since lying can be risky, caution may be beneficial. Once people catch you in serious contradictions, and once they start to share this interesting and relevant information with your friends, you will lose badly in the competition for prestige. You would have fared even worse in a society of a few hundred people where everyone know and depended on everyone else. Most of us simply don't have the skill to stray very far from the truth without getting caught. We can avoid danger by sticking reasonably close to reality. (Burling 2005: 196)

Filtering not only filters out false or misleading assertions, it also provides an incentive for speakers to not mislead in the first place.

Comprehension-with-filtering may not, as a matter of fact, be as reliable as perception. But for all that, it is still an awfully good way of acquiring information; it is a very reliable guide to the way things are. And the reliability enhancing aspect of filtering explains, at least in part, why we bother to filter what we take others to assert. Comprehension-with-filtering has inducing true beliefs reliably as a function. Getting informed is one thing, among others, that comprehension-with-filtering is for.

There are different kinds of warrant. Many comprehension-based beliefs are supported by justifications. But comprehension-based beliefs also enjoy entitlements. For beliefs based on normally functioning processes that have

forming true beliefs reliably as an etiological function enjoy entitlement, and comprehension-with-filtering has that function.²⁰

²⁰ This chapter was long in the writing, from conception to completion. I am grateful to Duncan Pritchard for his invitation to participate in the conference where some of the ideas in the chapter were presented in their infancy. At the conference I benefited from Finn Spicer's enthusiastic commentary and from questions from Sandy Goldberg, Alvin Goldman, Jon Kvanvig, and Ernest Sosa. A later presentation to a workshop at UCLA helped move it along. Thanks to Sven Bernecker, Tim Black, Peter Kung, Nikolaj Pederson, and Masahiro Yamada for their questions on that occasion. Discussions with Justin Coates and Patrick Todd spurred completion. The comments from a reader for Oxford University Press led to some last-minute improvements.