An Artifact is to Use

An Introduction to Instrumental Functions

1. Introduction

"Buttons are to keep people warm."
Ruth Krauss, *A Hole is to Dig*

One of the defining characteristics of artifacts is that they have functions; they are *for* something. Many, perhaps most, non-philosophical discussions of artifacts and their functions have a decidedly practical component.[1] In general, if one knows an artifact's function, she knows that the artifact can be used to realize certain goals. Knowing the function provides one with means to related ends she may wish to pursue. We interpret such means-end claims as suggested by Georg Henrik von Wright 1963: an end is a state of affairs one wants to realize and a means is an action—something to be done[2]—that can bring about the end.

Philosophical treatments of functions, however, have focused on the role of functional explanations rather than the instrumental role of functional knowledge. This interest is at least partly motivated by the emphasis in biological function that largely reintroduced functional terminology as respectable philosophy. We are interested in the functions of biological features in order to explain their presence, prevalence or persistence (as in (Wright, 1973; Millikan, 1989; Neander, 1991), et al.) or else to explain a capacity exhibited by a larger system under investigation (see especially (Cummins, 1975)). In either case, functions are of interest for the explanations they provide.

Functional explanations are also of interest in the case of artifacts. When we ask what the master brake cylinder is for, we are looking for an explanation: Why is this device under the hood of my car? How does it contribute to the vehicle's performance? The explanatory role of artifact functions arises also in archeology, in coming to understand why certain structures are commonly found in a historic community. Similarly, in reverse engineering, it is natural to analyze a system in terms of its functional components. In each of these cases, a theory of functional explanations captures the ways in which teleological concepts are being employed.

But artifact *users* are primarily interested in functions for practical rather than explanatory reasons. More often than not, when one considers the function of an artifact, he is interested in how it can be used to realize his ends—or ends he may adopt in the future. Artifact users are less often interested in teleological stories about how the artifact came to be where it is or why a particular artifact type is so common. Instead, they want to know how to *use* an artifact and engineers design artifacts to be *used*. Usage is inherently practical and functions are about how an artifact can be used, so we expect that functions have a clearly practical component. In particular, artifact use is goal-directed and so we expect a connection between function ascriptions and means-end relations.

The step from function ascriptions of the sorts found in, say, (Millikan, 1984) or (Cummins, 1975) to these practical consequences is not at all clear, however. To take an example, suppose that I know the Millikan-style function of a stapler. That is, staplers are for fastening papers together and this is their function either because this is what they were designed for or perhaps because this is what similar artifacts did well in the past. Their function is defined in terms of how the stapler came to be widely reproduced and found in certain settings. It is not so obvious how this information motivates my use of a stapler in appropriate situations. This very notion of function is directed to theoretical, not practical, knowledge. It is aimed at providing understanding about how things came to be as they are, which is a different task than giving practical options for achieving one's goals.

There is, nonetheless, a tempting argument which leads from knowledge of the stapler's function to practical consequences. It goes something like this: this stapler was manufactured because similar tokens have been useful in fastening papers together. Therefore, this token is also likely to be useful in fastening papers together and so if I need to perform this task, I have a reason to use this stapler. And perhaps this reasoning is correct, as far as it goes, but it does not seem a plausible reconstruction of how people in fact reason about artifacts. If I were asked what a stapler is for and I responded in terms of previous artifacts that performed a particular task which explains something like the reproductive success of the artifact type, I imagine that I would receive a rather odd stare.[3] In most questions of this sort, we are interested in direct, practical information and not a historical account from which this information can be derived.

A system analysis account like Cummins's fares little better. For Cummins, the function of the stapler is relative to a system and capacity under analysis. So, for example, the stapler is for fastening papers, because in a situation involving the stapler, a pile of papers and a user, the capacity for producing a stack of fastened papers depends on the role that the stapler provides. To be fair to Cummins, he was restricting his attention to the role of functions in *scientific* explanations and so this example is understandably awkward. But the primary point should be clear enough: analyzing a system is different than reasoning about what one should do and what tools will aid in a task. When I consider whether or not I should use the stapler, I am certainly *not* trying to explain a capacity in some system involving the stapler. Indeed, at this point, there is no system to be explained since the stapler sits unused. But even if we consider a hypothetical system in which I am using the stapler, it does not seem plausible that my deliberation involves the required analysis. Again, the focus on this sort of function is on theoretical rather than practical understanding.

Instead, we propose to introduce a new notion of function, namely *instrumental function*, aimed at clarifying the kind of functional knowledge that produces clear practical consequences. We will identify the characteristic features of instrumental functions and show how these features account for the means-end relations that both users and designers naturally associate with artifactual function. We make no claims that instrumental functions are the single “right” notion of function.

Indeed, Millikan- and Cummins-style functions succeed in their roles where instrumental functions would not: the latter concept is not well-suited for explaining either presence or capacity, just as the other notions of function seem disconnected from the practical role of functional knowledge.

There are nonetheless interesting relations between instrumental and explanatory functions. Instrumental function ascriptions can yield both Millikan- and Cummins-style function ascriptions—if we know an instrumental function, then we have reason to believe in related explanatory function ascriptions, as we will see. One may, therefore, be tempted to view instrumental functions as a kind of subclass of the two explanatory concepts, but we find this interpretation misleading. Instrumental functions reflect a practical fact about an artifact type, regardless of whether this fact figures in some explanation or not. That instrumental functions usually have explanatory consequences is an interesting side issue, not an essential feature.

Interpreting function ascriptions in terms of means-end relations is a fairly subtle task. Tokens inherit their functions from their types, but individual tokens may be in bad shape, unable to perform as they should. Thus, we introduce *normal tokens* in Section 4 as an essential feature of functional reasoning.

To be sure, our analysis may seem at times a bit removed from the actual practical reasoning in which folks engage, especially as it is understood in the Continental tradition of philosophy of technology. One's interpretation of technological artifacts depends on such contextual issues as a user's prior intentions, desires, beliefs, goals as well as the situation in which the artifact is found.[4] The presentation of instrumental functions offered here, on the other hand, omits these important contextual features of functional knowledge, since we aim primarily to contribute to the current discussion of functions in the analytic literature. If we can provide a sketch of instrumental function that works, at least in the abstract, to indicate the relation between technological function and practical reasoning better than the predominating alternatives, we will have accomplished our primary objective.

As a secondary objective, we will analyze artifact failure and malfunction in terms of the associated means-end relations. Malfunction has played an important role in the literature, including discussions in (Millikan, 1984; Neander, 1995; Davies, 2000; Schurz, 2001; Franssen, 2006; Vermaas and Houkes, 2003) and elsewhere, but the connection between malfunction and instrumental reasoning has been largely ignored.[5] We show that our account of function can provide a natural distinction between these two essential functional concepts.

2. The features of instrumental functions

In this section, we will examine the features of an instrumental function ascription. These four features, taken together, support the basic practical consequence of such ascriptions—very roughly that, if t is for bringing about φ, then *using* t is a means to the end φ.

We argue that, in order to reliably infer these consequences, one must be able to answer four basic questions about the instrumental function in question. These questions are:

(a) What is the aim of the function?

(b) How should a token be used to accomplish the aim?

(c) When and where should a token be used for this end?

(d) What function-bearing type are you talking about?

The answers to these questions correspond to the following features.

(A) a functional goal,

(B) a use plan,

(C) a set of normal contexts of use.

(D) an artifact type.

We regard these four features as characteristic of instrumental functions, but they also correspond roughly to the features of Cummins systems described in (Millikan, 2002, p.120). The artifactual type corresponds to the Cummins system and the functional goal to the output capacity being analyzed. The use plan is roughly analogous to the "allowable inputs" to the system and the contexts of use appears as Millikan's "allowable conditions of operations". Our account differs from her discussion by placing emphasis on the level of artifact types and their attributed functions rather than on the analysis of system capacities. Again, this difference is crucial to our aims: practical deliberation involves reasoning about how to use the things at hand, not analyzing the capacities of an existing system.

Artifacts are not the only bearers of instrumental functions. Non-artifacts, including biological features, may also have instrumental functions. A dairy farmer knows the instrumental function of a cow's udder. It is for getting milk. There is a way in which the udder can be manipulated to achieve this goal, at least in the right circumstances. Of course, most biological features—even those with clear explanatory functions—do not have any evident instrumental functions. One may know what the frog's heart is for in, say, Cummins's sense, but this does not entail any clear practical consequences at all. In any case, we are particularly interested in artifacts and their instrumental function and our presentation will be similarly biased toward artifactual functions.

We will discuss each of the features (A)–(D) in turn.

Functional goals

Every instrumental function includes a *functional goal*: a condition which can be realized by properly use of the item. A function ascription asserts that the artifact is good for *something*, specifically, for bringing about some particular state of affairs. We use φ to denote functional goals, since such states of affair are commonly expressed as propositional functions. But we do not assume that such goals are Boolean: some goals (such as, "drying hair in a timely manner") can be satisfied to greater or lesser extent. We also do not assume that the user can reliably evaluate the degree to which the goal has been satisfied, but we *do* assume that there is a fact of the matter involved.

It is worth emphasizing here that we view functional goals as *propositional functions* and not simple sentences.[6] In logical terms, they include (typed) free variables: staplers are not for fastening this or that stack of papers, but for fastening stacks of papers. This functional goal may be crudely represented as $\mathbf{Fasten}(x)$, where x is a variable ranging over stacks of papers of appropriate thickness. A particular application of the stapler will involve a particular stack s of papers and will be successful just in case, after the use, $\mathbf{Fasten}(s)$ is true.

In this respect, our account is similar to the discussion of *derived proper functions* in (Millikan, 1984). In Millikan's terms, if $\mathbf{Fasten}(x)$ is the direct proper function of staplers and s is a stack, then $\mathbf{Fasten}(s)$ is a derived proper function. The terminology seems a bit awkward—are staplers really for fastening s, even derivatively? Normally, functions are stated in more general terms than this. Our contexts of use will play a similar role, but without Millikan's terminology.

Use plans

The connection between instrumental function and the related goal is so intimate that we colloquially confuse the two. We say that "Staplers are for fastening papers together" is a function ascription, but there is nothing explicit in that sentence aside from an artifact type (stapler) and a functional goal (the fastening together of papers). Nonetheless, if this *is* an instrumental function ascription, it must include certain other implicit claims. In particular, one can reasonably infer that there is a standard way of *using* a stapler for fastening papers together. It is surely not the case that one requires a novel, *ad hoc* procedure to fasten each pile of papers. function ascriptions are about particular ways of using an artifact.

Thus, every instrumental function explicitly or implicitly includes a *use plan* (Houkes, 2006; Vermaas and Houkes, 2006), which we denote α. Credible claims about artifact capabilities assume prescriptions for how the artifact should be used to realize the functional goal. Such prescriptions are constructed or discovered during function creation. It would be absurd to claim that one has designed an artifact for a particular function but does not know how the artifact should be used to do this task. Similarly, discovery of an accidental function must involve some plan for how to realize the functional goal. For a more thorough discussion along these lines, see (Houkes, 2006).

A use plan is a prescription for how one should manipulate the artifact and related objects in the context of use in order to realize the goal. We allow for broad differences in the expression of such plans. They may be explicit and detailed ("Tighten the filter one quarter turn after initial contact.") or vague and broad ("Always obey local traffic laws."). They may include conditional actions ("If the stapler is empty, load it."). But in each case, they describe what one should *do*. Thus, as we will see, use plans provide the *means* for our means-end analysis.

One may object that some artifacts passively realize their goals and so it is unclear whether they have associated use plans. For instance, components in a complex system (such as the master brake cylinder in a car) do not seem to come with user prescriptions. Similarly, how does one "use" a retaining wall? In both cases, one may argue that there is no *use* evident. The artifact just does what it is supposed to in the right circumstances, much like the human heart.

On the contrary, the user does indeed have something to do in each of these cases. He has to ensure the cylinder or wall is installed properly and thereafter he should perform regular inspections and maintenance. This is the use plan for these artifacts. A retaining wall will not prevent erosion unless it is installed properly and is in good working condition. If one wants to prevent erosion, then she must make sure that the wall is installed and maintained (perhaps by delegating the task to a competent party). Similarly, the driver should be sure that a qualified mechanic installs and maintains the brake cylinder as part of the general maintenance of the vehicle. Unless she does this, she cannot reasonably expect the end she desires, namely, that the brakes slow and stop the car when needed.

Such use plans seem superficially similar to prescriptions that heart-bearers should follow ("Schedule regular medical checkups."), although it does not seem as if the heart's function is instrumental. In fact, the connection between use plans and instrumental functions is stronger than the non-instrumental analogues. The development of such plans is a necessary condition for deriving practical consequences, but the discovery of explanatory functions does not involve any analogous plans in general. The scientist who discovers the function of a frog's brain may not know how to keep it in good working order, but the engineer who designs brake cylinders must be able to say how they should be installed and what physical features must be maintained to ensure reliable performance.

In our account of instrumental functions, we clearly distinguish use plans from functional goal, while in other function theories, the two seem to be conflated. For instance, Larry Wright defines "Z is a function of X" as the conjunction:

1. X is there because it does Z.

2. Z is a consequence (or result) of X's being there.

In this definition, it is not clear what type Z has. In (1), it appears that Z is an activity, something which is *done*. In (2), on the other hand, Z is evidently a condition or state of affairs, the sort of thing that can be a consequence of another state of affairs (X's being there). There seems to be a type confusion here. This confusion may be forgivable if our goal is to explain the presence of the item, but it is essential that we distinguish ends from means if we are to derive the instrumental relations necessary for practical reasoning.

Loosely, then, things which are *done* are part of the use plan while *end states* toward which the action aims are part of the functional goal. Unfortunately, natural language expressions can confuse the distinction. For example, "Staplers are for fastening papers together," appears to identify an activity ("fastening papers together") as the functional goal, but this is just a misleading oddity of language. One does not use staplers because he desires to experience the *act* of fastening papers, but rather because he wants the end result, namely, that the papers are fastened.[7]

Contexts of use

Instrumental functions include prescriptions regarding not only *how* but *when* one should use an artifact. That is, such functions include a *specification of contexts* in which the artifact can be used to realize the functional goal. One cannot make reliable practical decisions about artifact use unless he can identify situations in which use is appropriate, i.e. likely to result in a suitable outcome. These specifications include descriptions of users and other objects involved. For instance, staplers are applied to small stacks of papers and cars are operated by persons who know how to drive.

Conceptually, a specification of contexts of use amounts to identifying a set C of situations in which the artifact may be used. We write $c \in C$ to indicate that c is a situation satisfying the specification C.

Specification of contexts serve three distinct roles.

(i) They limit the situations in which an artifact is expected to perform its function. A car should not be expected to provide reliable transportation if its operator does not know how to drive (does not have *operational knowledge*, in the terminology of (Houkes, 2006)).

(ii) They provide parameters for the use plan and functional goal. One uses a stapler by inserting a stack of papers and pressing down on the mechanism. Like functional goals, use plans include free variables: $\mathbf{Insert}(x); \mathbf{Press}$ where x is a variable ranging over stacks of papers. A particular usage includes choosing an appropriate stack s and executing the use plan with s. Thus, contexts serve as a bridge from the general (use plans and functional goals) to the specific (applications of use plans and evaluation of goals in context).

(iii) Success can be context dependent. Brakes should stop cars on both wet and dry pavement, but we expect shorter stopping distances on dry pavement.

Like the other features we've identified, the contexts of use can evolve over time. It is unlikely that early designers of ballpoint pens explicitly considered the presence of a gravitational field as part of the normal context of use. Nonetheless, the fact that such pens don't work in zero gravity is relevant at least to NASA, if not to the rest of us.

Contexts of use are missing in some function theories while implicitly present, at least in part, in others. They play no obvious role in (Vermaas and Houkes, 2006), for instance. On the other hand, Ruth Garrett Millikan's *Normal conditions* 1984 seem very similar to our specification of contexts, at least regarding role (i), and her derived proper functions play roles similar to (ii) and (iii). Also, user abilities and circumstances are an explicit part of reasons to use an artifact in (Franssen, 2006), essentially role (i).

Artifactual types

Lastly, instrumental functions apply to specific *types* of objects (typically, artifact types), denoted T. This may be a matter of some controversy, since some accidental functions appear to apply only to particular tokens rather than types. For instance, when one needs to retrieve his keys from a grating into which they've fallen, he may apply some chewing gum to the end of a branch and try to pull them up by adhering the sticky gum to the keys. It may seem implausible that this device is part of a larger artifactual type, even though it is a well-known solution to the problem.

Token-level instrumental functions are an interesting topic, but we save it for later work. The restriction to type-level functions allows a richer development of our theory. In particular, Section 5 deals with malfunction and this concept relies on a sense of normal tokens of similar type. In the rare case that a token is truly novel and not an instance of a larger type, concepts like malfunction may not apply, since malfunction is (as we will argue) a comparative term.

In fact, it seems plausible that instrumental functions *do* primarily refer to artifact types and only derivatively to tokens. One may argue that, even for novel artifacts, the proper subject of a function ascription is an artifact type, albeit a type instantiated by a single token. Some realists may object to the proliferation of types, however, and we prefer to avoid such ontological distractions. So let us restrict our attention to functions that *do* apply to types and avoid this controversy.

Artifact types can be broad or narrow. Because our function ascriptions include use plans, we require that our functional types are narrow enough so that a common use plan can apply to each token. A broad type like “bottle opener” can be realized in many different ways and with many different use plans, and so is too broad for our needs here. Instead, means-end relations are introduced by narrower subtypes like “corkscrew”.

These four features—functional goal, use plan, contexts of use and artifact type—are products of the design and discovery processes that yield instrumental functions. One cannot plausibly claim the he has discovered an instrumental function or created an artifact with such function unless he can identify these four features: productive usage requires knowing how and when to use the item, what should be achieved and what class of items are under consideration. Admittedly, it is sometimes hard to discern each of these features in informal talk about functions. We often speak tersely about functions, giving only the type and the functional goal. But such terse ascriptions cannot generate clear means-end relations unless the missing features of use plan and context specification are assumed to be implicit.

Moreover, an instrumental function ascription can be characterized by these four features. That is, any two ascriptions which involve the same features φ, α, C and T are identical—they both express the same instrumental function. And two ascriptions with different features express different functional assertions. Thus, we will identify a function ascription by the tuple $\langle\varphi, \alpha, C, T\rangle$.

As we have said, if one believes an instrumental function ascription $\langle\varphi, \alpha, C, T\rangle$, then he should also accept some related practical consequences. Indeed, that the ascription is true means, in part, that T-tokens can be used to realize φ; if T-tokens could not be so used, then it would make little sense to say that is what they are *for*, i.e. that this is an instrumental function. Consequently, let us offer the following tentative principle, to be revised later.

The instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$ entails that, in situations satisfying C, using a T-token as prescribed by α is a means to φ. (**ME-1**)

Note that the means-end relation here is that of a sufficient means to an end. Functions do not provide necessary means, but instead suggest one way to realize the functional goal.

We can see now the relation between instrumental functions and Cummins-style functions. The instrumental function $\langle \varphi, \alpha, C, T \rangle$ entails that, in situations satisfying C in which a user is using a T-token appropriately, he can bring about φ. Thus, the system consisting of the user, the T-token and whatever other elements of C are relevant has the capacity to bring about φ and this capacity is due to the causal contributions of the T-token. Hence, an instrumental function plays a natural role in a Cummins-style analysis of relevant systems.

Since we can apply (**ME-1**) for each situation satisfying C and for each T-token, a function ascription entails a *family* of individual means-end relations: for each $c \in C$ and $t \in T$, using t in situation c as prescribed by α is a means to φ. Indeed, this claim is implausible on the face of it: some T-tokens may be broken and hence unable to serve their purpose. In those cases, we do *not* expect the means-end relation to hold. We will return to this caveat in Section 4, but for now, let us take (**ME-1**) as stated above.

One may well be mistaken about the efficacy of T-tokens in achieving φ. Perhaps T-tokens *cannot* be used to bring about φ after all. But functional knowledge is no different than practical knowledge in general in this respect. One may always be mistaken about whether a particular action really is a means to a given end. The principle expressed in (**ME-1**) still stands (with the caveats discussed hereafter): one is justified in believing the function ascription only to the degree that he is justified in believing the associated family of means-end statements.

To clarify, some means are more reliable or effective than others. For all its faults, a snowblower is a more effective means of clearing a sidewalk than a snow shovel. This does not entail, however, that one is more justified in the belief that a snowblower is for removing snow than in the belief that a snow shovel is also for removing snow. The principle (**ME-1**) does not depend on *efficacy*. As long as both types (snowblower and snow shovel) provide means to the same end, each has the function of removing snow (albeit with different user plans and context specifications).

3. The teleological nature of instrumental functions

The family of means-end relations in (**ME-1**) is a necessary condition for the instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$ to be true. It is not sufficient, since there is more to instrumental functions than mere capacities. Modern automobiles, for instance, can be used to contribute greenhouse gases to the environment. There is a clear use plan and set of contexts in which they are effective at doing so. Nonetheless, this is not an instrumental function of automobiles: they *can* be used this way, but this is not what they are *for*.

The distinction we want is that between capability and purpose,[8] but the purpose here is not so problematic as some other theories of function face. For explanatory theories that include functions of frogs' brains and bats' wings, the teleological nature of function is fairly delicate. Frogs' brains serve a purpose, but they do not serve *someone's* purpose—the teleological terminology does not refer to the intentions or goals of a rational agent. It must either be explained away or explained in purely naturalist terms. As (McLaughlin, 2001) says, "A functional explanation [of the sort found in science, say] is independent of any attribution of intentionality."[9]

But we have no qualms about using intensional terms in our definition of instrumental functions. After all, these are the functions that *agents* reason about. Our interest is in deriving the means-end relations that play a central role in instrumentalist accounts of practical reasoning. Desires, intentions, goals, etc., play an equally important role in such accounts and so it is perfectly appropriate to use similar terminology in our definition of instrumental function. The function of an automobile is transportation, and not making greenhouse gases, because people want or need the former capability but not the latter.

More to the point, people are interested in autos because they are capable of transporting persons. Autos are valued, at least in part, for this reason, while they are not valued for their capacity to create greenhouse gases. We produce, purchase and maintain autos just because they provide a means of transportation. Thus, we suggest that the difference between *function* and *(mere) capability* is that if a type has a function $\langle \varphi, \alpha, C, T \rangle$, then we value tokens of that type for the fact that they can be used as a means to φ in the sense given in (**ME-1**).

We may be more explicit about what it means to value tokens of an appropriate type for the sake of an instrumental function. For this, we will adapt a principle first presented by (Sorabji, 1964) and later developed (under the name "Sorabji's Rule") by (McLaughlin, 2001). Sorabji distinguished two senses of "function". His functions of sense (i) serve as a rough analog of instrumental functions. For a type to have a function of sense (i), it is necessary that we are willing to expend some effort so that at least one token of this type has the proper capability. Automobiles may have the function of transportation because we have shown our willingness to expend effort to design, produce, procure, fuel and maintain them so that they are capable of producing the right effect. But this willingness need not result in actual expended effort: a fallen log may have the function of being a bridge just so long as I am ready to put in some work *should it become necessary* so that the log continues to serve this purpose.[10] If, as it happens, the log continues to serve this purpose with no intervention on my or others' parts, it is nonetheless true that the log functions as a bridge, by virtue of the fact that we have adopted this positive attitude towards the capability.

The effort I am willing to expend must furthermore be causally relevant to the artifact's function. As McLaughlin says, it is not enough that we appreciate the collision of two continental plates in order to confer the function of skiing location upon the Alps. There is no effort which we could expend that would effect or encourage this collision. For an agent to be capable of conferring a function on an item, there must be some action which he *could* perform that would preserve the item's relevant capacity. In other words, the function-conferring agent must be *causally relevant* to the function-bearing item.

Thus, we have identified two necessary conditions for instrumental function: first, that tokens of the type can be used as a means for the functional goal in appropriate situations—i.e. the principle labeled (**ME-1**)—and second, that tokens of this type are valued for their capability to be used thus. These two conditions must attain if we are to correctly claim that this type has the ascribed function. Moreover, we claim that these two conditions are jointly sufficient to justify an instrumental function claim: if tokens of type T can be used as a means to φ as in (**ME-1**) and if such tokens are valued for this reason, then it is a function of T-tokens to realize φ (by doing α in situations satisfying C). This is what T-tokens are for:[11] this is how they can be used and this is a reason people are interested in them.

We therefore offer the following preliminary definition of instrumental function (to be revised in Section 4).

> *The following are necessary and jointly sufficient conditions for an instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$ to be true.*
>
> *(a) In situations satisfying C, using a T-token as prescribed by α is a means to φ.* (**IF-1**)
>
> *(b) Some causally-relevant persons value the capacity of T-tokens described above.*

Thus, an agent who accepts the function ascription $\langle \varphi, \alpha, C, T \rangle$ is also committed to the claim that someone (not necessarily him) values T-tokens for this capacity. Similarly, if he accepts that T-tokens can be used as described and that someone values T-tokens for this reason, he should also accept the function ascription $\langle \varphi, \alpha, C, T \rangle$.

We can now describe how instrumental functions typically produce Millikan-style functions. If tokens of type T are valued by their capability to bring about an end, then we have a reason to collect, maintain and reproduce tokens of this type. Thus, the utility in bringing about

φ can often explain the prevalence of T-tokens and hence this capacity is a proper function in Millikan's sense.

4. Normal tokens

The conditions offered thus far have ignored a basic fact about function-bearing artifacts, namely that individual tokens of an artifact type can differ widely in their capacities. A poorly maintained bicycle will not provide transportation as reliably as a well-maintained bike. Indeed, it may become incapable of providing transportation at all. But (**ME-1**) states that using a bicycle appropriately is a means for moving from one place to another. It is natural to read this as: using *any* bicycle would effect our goal. Let us call this the *universal interpretation* of (**ME-1**). It appears that the universal interpretation is incorrect. Only *some* bikes are capable of satisfying our needs.

We may defend the universal interpretation of (**ME-1**) by taking our context specification C to be sufficiently narrow and include conditions on the token t. That is, we could require that C includes a description of what a "normally functioning" T-token is. Since (**ME-1**) applies only for contexts satisfying the specification C, our expectations apply only to those tokens that meet certain structural requirements. This approach is again similar to Millikan's Normal conditions [1984], since the condition of the artifact plays a role in the Normal explanation of its function.

We avoid this alternative for two primary reasons. First, with complicated artifacts, the user is unlikely to know what structural features are relevant for use or whether a particular token possesses such features. An engineer may know how the wires inside my television should be connected, but I surely do not know this and I am unable to easily confirm that they are properly connected in any case. So this precondition for using my television cannot be part of my understanding of its function. The second (and more significant) reason for keeping a notion of normal token separate from conditions of use is discussed in Section 5 when we give a preliminary analysis of malfunction. For this, we must distinguish environmental abnormalities in applications from abnormalities in the token used.

Principle (**ME-1**) stated that an instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$ entails that T-tokens, when used appropriately, are capable of realizing φ. Indeed, there is something more to instrumental functions than this: T-tokens when used in the appropriate way and in the appropriate setting should bring about φ in some particular way. Bikes serve as transportation by transmitting the force applied to the pedals to the wheels and propelling the bike forward. For a particular bike to be capable of performing in this manner, it must have certain relevant physical features, including pedals that can withstand the force applied, a chain capable of transferring the force to the rear wheel and two wheels that can balance and propel the bike. Without these features, we cannot expect that the bike will serve as a means of transportation.

A bike is in *normal condition* (i.e. is a *normal token* of the type bicycle) just in case it has the physical features necessary to be capable of realizing its functional goal in the appropriate setting. In general, a normal token of type T is one with all of the requisite features necessary to realize each of the functional goals of T in the manner intended.[12] With this definition of "normal token", the following principle appears to be a triviality.

> *The instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$ entails that, in situations satisfying C, using a* normal *T-token as prescribed by α is a means to φ.* (**ME-2**)

Nonetheless, we do not regard (**ME-2**) as a useless tautology. First, we considerably strengthen its claim by adding:

> *The instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$ entails that normal T-tokens are physically possible constructions.* (**Norm**)

Principle (**Norm**) thus adds the restriction that, even if there are no extant T-tokens capable of performing the function ascribed, such tokens can be constructed. Indeed, if such tokens were not possible, it would be hard to explain how T-tokens came to be valued in accordance with (**IF-1**).

More importantly, our notion of instrumental function is intended to fit into a practical account of how epistemically-limited rational agents use artifacts. Such agents do not typically know all of the features relevant for the adequate performance of an artifact. Instead, their notions of normality are a product of experience and testimony and their ability to judge whether a particular token is capable of performing as required is fallible. Thus, the intended meaning of (**ME-2**) can be restated thus:

An agent who accepts the instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$ and accepts that t is a normal T-token also accepts that, in situations satisfying C, using t as prescribed by α is a means to φ. (**ME-2′**)

Consequently, in such situations, if the agent intends to bring about φ, he has a reason to use t.

In many respects, an agent's limited grasp of normality is similar to his limited grasp of the user plan and context specification. Each of these features develops over time, either through personal experience or testimony from others. Indeed, even the state-of-the-art engineering knowledge for each of these features of instrumental functions are prone to change. Let us take an illustrative example.

Suppose that an agent has never before used a television and he has acquired knowledge of its function, either through the user manual or informal communication. He is told that the power button turns the set on or off. Since he has some knowledge about buttons on electronic devices, he concludes that the power button must be present and in good working condition. Of course, he is unable to see whether the button really is in good condition, since most of the apparatus is inside the set. Nonetheless, in at least some cases, he can see that the power button is not in an appropriate condition (it is broken off, say, and hard to press) and thus that the set is not a normal token in at least this respect.

If he has some acquaintance with electronic devices and he sees that the television has a power cord, he can immediately deduce that this cord is relevant for the television's function. His concept of normal television therefore includes a power cord in good condition.[13] A cracked picture tube will certainly be diagnosed as an abnormality, especially after he uses a set and sees how the picture tube contributes to the function. Working artifacts tend to have a regular appearance, so that even if our agent does not know how a television works, he may be able to distinguish an antenna in good condition from one which has been broken off, leaving a jagged crimp at the end. Because he will likely suppose that the antennas are there for a reason, he will likely consider the condition of the antennas as relevant for the functioning of the television. As his experience with sets grows, his notion of the features relevant for functioning becomes more detailed and sophisticated, so that he can more accurately predict whether a set will perform as expected before use.

The stronger the association between a feature and the manner in which an artifact functions, the more decisive the feature is for judging that a token is normal. A picture tube is evidently essential for the proper functioning of a television and so a set with a cracked or missing tube is certain to fail to function as it should. The contribution of an antenna is less obvious to our novice user and thus he will be less certain whether the set will work with a damaged antenna. For his purposes, a set with a damaged antenna is closer to a normal condition (i.e., more likely to perform as it should) than a set with a cracked tube, although neither fit his concept of a perfectly normal token of the type.

The electrical engineer has more detailed and precise knowledge regarding televisions than the typical user. He understands better how a television is supposed to realize its function and so can develop a more sophisticated account of the relation between physical features and artifact behavior. But the difference here is largely in the degree and completeness of the account; both user and engineer come to their concepts of normality via their understanding of how an artifact type is supposed to work. The engineer may rely more on knowledge regarding design, physics, and so on, while the user depends primarily on naive induction from previous uses, communicated use plans etc.[14] The more detailed understanding of the engineer leads to better predictions regarding artifact behavior and thus more reliable beliefs regarding the corresponding means-end relations.

A type may be badly designed, so that no tokens of that type are capable of performing their intended (or advertised) function. For instance, Beth Preston suggests in [1998] that bug zappers are not really capable of reducing the number of mosquitoes in the area. According to Preston, then, bug zappers cannot have the function of reducing the number of mosquitoes, since this is something they cannot do in the settings in which they are supposed to be used. Instead, zappers must have the function of (merely) giving the *appearance* of reducing the number of bugs. This analysis is consistent with our principle (**ME-2′**). One may buy (or even design) a bug zapper in order to reduce the mosquito population, but in this case, one is drawing a practical inference from an incorrect premise, namely that bug zappers serve this function.

Our beliefs about normal tokens may be mistaken in other ways—we may be wrong about what structural features explain relevant behavior or whether an actual token exhibits these features. There are many judgments that can go wrong in our step from function ascription to practical consequence, but that is a feature of reasoning generally and practical reasoning in particular. It is tough to get the premises right, but the principle in (**ME-2′**) seems nonetheless correct.

With these considerations in mind, we offer the following revision to (**IF-1**).

> *The following are necessary and jointly sufficient conditions for an instrumental function ascription* $\langle \varphi, \alpha, C, T \rangle$ *to be true.*
>
> *(a) In situations satisfying* C*, using a* normal T*-token as prescribed by* α *is a means to* φ*.* (**IF-2**)
>
> *(b) Such normal* T*-tokens are physically possible constructions.*
>
> *(c) Some causally-relevant persons value the capacity of* T*-tokens as described in Section 3.*

An agent who accepts an instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$ is therefore committed to principles (a)–(c), where the term "normal" corresponds to the agent's concept of normality for T-tokens. As well, an agent who accepts (a)–(c) should also accept the corresponding instrumental function ascription.

There is clearly more to say about the use of normal tokens in means-end relations derived from function ascriptions. The sketch we give here is preliminary but sufficient for this introduction to instrumental functions. In the following section, we put the concept of normality to work.

5. Failure and malfunction

Malfunction serves as a kind of litmus test in (Millikan, 1989). She writes, "an obvious fact about function categories is that their members can always be defective... hence unable to perform the very functions by which they get their names." Etiological theories can account for malfunction while system capacity theories like (Cummins, 1975) cannot,[15] and hence (according to Millikan) the former are preferable for this reason. The importance of malfunction is echoed in (Neander, 1995; Schurz, 2001; Vermaas and Houkes, 2003) and elsewhere.

Certainly, a theory of instrumental functions must include some notion of malfunction. An agent's plan depends on the working state of the artifacts he has at hand. It would be foolish to use an bicycle with a broken chain as a means of transportation and so we should take care that our theory of instrumental functions includes the distinction between well- and malfunctioning artifacts.

The literature is surprisingly superficial on the actual meaning of "malfunction", however. A brief definition is suggested in Millikan's above quote and also in Preston's [1998] paraphrase of Millikan: "If you can say what a thing is supposed to do, then you can also say when it is failing to do something that it is supposed to do, that is, malfunctioning." In (Neander, 1995), we find that "a biological part functions properly when it can do what it was selected for and malfunctions when it cannot."[16] Presumably, the case is similar for artifacts.[17]

We see two features common to each of these accounts. First, malfunction is about capability. It is not about whether the artifact did what it was supposed to on some particular occasion, but what it would do *if* used. Second, malfunction applies to individual tokens, not types. Types may be poorly designed, so poorly that—like Preston's bug zapper—no token can do realize its functional goal, but types do not malfunction.[18]

Failure is similar to malfunction in this respect: it applies to tokens rather than types. But failure, unlike malfunction, is about individual performances. An artifact token t either *fulfills* or *fails to fulfill* its function in a particular application. This judgment applies to a specific instance in which the token was used according to use plan α in some situation realizing a context in C. We suggest the following definition:

> A token t *fulfills its function* in a particular application just in case the functional goal φ was realized as a result of that application. Otherwise, it *fails* to do so. (**Fail**)

Fulfillment and failure may be a matter of degree, since φ may be realized to greater or lesser degree.

Thus, a token t fulfills its function if the goal is realized *due* to the application of the use plan with respect to t. Thus, the application must be causally relevant to the realization of φ, but even this requirement is perhaps a bit liberal. In practice, we may also require that the application realizes the goal *in the intended way*, rather than by accident. A varmint rifle that misfires, striking a branch that falls on and kills the rabbit at which one was aiming, has failed to perform as intended even though the shooter's goal was realized. Let us ignore that complication at present and take (**Fail**) as our working definition of fulfillment and failure.

While failure is about a particular application, malfunction is a broader claim. A well-functioning token may fail to realize its goal; a perfectly good anti-aircraft missile may miss its target, for instance. Thus, failure is a negative evaluation of a particular outcome, but it does not mean that the token used was somehow bad. A malfunctioning token, by contrast, is one that cannot be expected to reliably or effectively bring about the functional goal. Malfunction is about hypothetical applications, what would happen if the token were used.

There are other senses of malfunction which we do not address here. As Sven Ove Hansson has pointed out (private correspondence), some malfunctions involve unintended consequences rather than inability to bring about an intended outcome. A television that presents a crisp image while filling the room with deadly radiation is surely malfunctioning, even though it is realizing its primary function. This kind of malfunction has been largely ignored in the literature so far, (Franssen, 2006) notwithstanding. We will focus on malfunction-as-inability for now, since it is more closely related to means-end relations than malfunction-as-undesirable-propensities.

The concept of ability (or, better, *capability*) needs some attention, however. Suppose one has a pistol with a weak spring in its firing mechanism. Most times the trigger is pulled, the hammer does not strike hard enough to fire the primer, but sometimes the gun fires as desired. The pistol is therefore *capable* of doing what it is supposed to, but it does not do so *reliably*. It seems clear that this pistol is malfunctioning and hence that reliability is a consideration relevant for malfunction.

Perhaps we can tweak our functional goal to include reliability. Maybe the functional goal of a pistol is to discharge a bullet every time (or almost every time) that the trigger is pulled. But this fix comes with a price: we would lose our definition of failure. Failure is about particular applications, but "discharge a bullet every time" quantifies over a range of applications. It makes no sense to ask whether a particular application realized the goal "discharge a bullet every time the trigger is pulled."

Instead, we propose putting the concepts of reliability and effectiveness into the definition of malfunction.

> A token t is *malfunctioning* with respect to a proper function if it is unable to *reliably* or *effectively* realize φ in some situations c satisfying C when used according to α, i.e. if $\alpha(t)$ is not a reliable or effective means to φ in such situations. (**Mal**)

Thus, like Millikan, Neander, et al., we define malfunction in terms of capability, but with additional considerations of reliability and effectiveness.

Our definition (**Mal**) contains two undefined terms, namely "reliably" and "effectively". The concepts of reliability and effectiveness are inherited from the semantics of means-end relations. A means to an end may be more or less likely to realize its end. That is, in a particular situation c, there is some probability that doing an action α will realize φ. We call this probability the *reliability* of α as a means to φ in c. *Effectiveness* is the degree to which the functional goal φ would be realized as a result of α. Of course, the situation can be considerably more complicated than this, since a means may have low probability of very effectively realizing its end and a high probability of a less effective outcome. The two measures are thus not orthogonal in practice, but let us ignore these complications for this initial sketch of malfunction.

Definition (**Mal**) requires a comparative notion: a token malfunctions if it is not reliable or effective, but compared to what? Here again, we rely on our intuitions regarding normal tokens. A token is malfunctioning if it is considerably less reliable or effective at fulfilling its function *than normal tokens of the same type*. Thus, malfunction involves a comparison between the behavior of an actual token and the behavior one expects of normal tokens.

As a consequence, a malfunctioning token is not a normal token, since it does not behave as well as a normal token. But this leads to an obvious difficulty in assessing malfunction. Suppose that our television set appears normal in every way—that is, it exhibits every feature we associate with normal tokens—but that pressing the power button fails to turn the set on. This failure requires some explanation, since normal power buttons do not fail to work. We may wonder whether our use plan is correct. Perhaps we were mistaken in how we are supposed to manipulate the set. In this case, however, we have considerable confidence that pressing the button should activate the set. Perhaps we are mistaken about whether the current situation satisfies the context specification: the set may be unplugged or the outlet may not be powered. It is also possible that we may have to revise our understanding of the context specification, but again, in this simple case, that seems unlikely. In certain situations, we may be tempted to revise our belief in the instrumental function itself. Maybe we were simply wrong that this artifact type does what we thought, but again this is unlikely in the current example.

Most likely, we would check to see whether the failure was an aberration. We would try pressing the button a few more times. If the set still failed to turn on, we would be forced to admit that the set is malfunctioning. It is not a normal token, contrary to appearances. We would investigate further the manner in which it fails to act as a normal set, most likely by taking the set to a repairman who has a more detailed grasp of the features of a normal television and the capability of checking which of these features is lacking in our malfunctioning set. Thus, repeated failures serve to suggest that our set is not performing reliably and hence is not normal, despite our initial assessment.

There is a further complication to our account of malfunction via comparison with normal tokens. The notion of normality is relative to an artifactual type, and types have a hierarchical structure. A particular phone may be a token of its model type, of wireless phones, of push-button phones, and of telephones *simpliciter*. A user will typically have different beliefs about the normal tokens of each of these types, and this may yield conflicting judgments about whether the token is malfunctioning. If the wireless phone loses connection with its base outside of a short radius, then it may be behaving less effectively than normal wireless phones. But perhaps early models of wireless phones had shorter ranges than later models. If this token is an early model, then it is behaving as designed. Should one say that this phone is malfunctioning?

Claims regarding malfunction should involve normal tokens of a *suitably narrow type*. "Suitable narrowness" is unfortunately difficult to specify precisely. We often compare our old wireless phones with wireless phones of similar age, for instance, but should we restrict our attention to phones of the same model? What criteria determine suitable narrowness for this purpose? Suppose that telephones made on Monday mornings are typically less reliable than telephones made at other times during the week. Should we compare our Monday morning phone to "normal" Monday morning phones before declaring that our token is malfunctioning?

We should not expect a rigorous definition of "suitable narrowness". Malfunction claims are context-sensitive and so we may expect that the relevant artifactual type from which we draw our normal tokens is similarly context-sensitive. Consider a token t of a particular model of stapler, say Acme. Suppose that this model is not particularly reliable: it often mangles staples rather than closing their tines, and so fails to fasten papers in situations in which better designed staplers work well. Our token t may behave as well as expected for Acme-tokens, but more poorly than normal staplers. In many situations, we would say that t is malfunctioning, and in other situations, that t is functioning as expected but is poorly designed. We would like our semantics for malfunction to accommodate both judgments: to allow that in some situations, normal behavior for the type "stapler" is relevant and in others, normal Acme-behavior is what matters to us.

Let us be satisfied with loose guidelines, then, that allow one to exclude too-broad and too-narrow artifactual types from consideration. To restrict breadth, we will use the same criteria we discussed in Section 2. The artifactual type must be relevant for the function discussed and so it must be narrow enough that a single user plan suffices for every token of that type. For narrowness, we suggest the following: the narrowest "suitably narrow" type is the type defined by the *design* of the token at hand. That is, the token at hand is the result of some design or selection process and that process defines a relevant artifactual type. Thus, for typical manufactured goods, we will sometimes consider normal tokens of the same *model* in order to make malfunction judgments, but we will not consider narrower types like "Monday morning telephones". We may sometimes want to distinguish bad design from malfunction and for this, we need design-defined artifactual types. But other historical features are not relevant to malfunction judgments. That a particular token was badly manufactured, for instance, explains rather than excuses its malfunctioning.

Thus, our constructed means-end relations provide a natural analysis of failure and malfunction. The latter also involves a comparison between an actual token and normal tokens of the same (suitably narrow) type. This comparative character has been largely omitted in the literature, because the relation between malfunction and the twin concepts of reliability and effectiveness has been overlooked. But reliability and effectiveness are a natural feature of our concept of instrumental functions and their practical consequences.

Notes

[1] Maarten Franssen even suggests in 2006 that evaluations of artifacts ("This is a good/poor screwdriver.") are really judgments about instrumental utility.

[2] Sometimes, one speaks of an object as a means, as when we say that a bridge is a means to crossing a river. We regard such locutions as shorthand for some action involving the object: in this case, walking across the bridge is the means we have in mind. Regardless of whether every object-as-means can be reduced in this way, we are interested in means-end relations in von Wright's sense here.

[3] And, indeed, Millikan is not claiming that her definition of "proper function" is an account of how individuals reason about artifacts. She is explicit in 1989 that this is a *theoretical definition* and not an exercise of conceptual analysis.

[4] Thanks to an anonymous reviewer for pointing out the importance of psychological context and to Krist Vaesen for his insights regarding environmental context.

[5] (Franssen, 2006) is a notable exception, although he gives a reason-based definition of malfunction rather different than our capacity-based analysis.

[6] Millikan takes a similar approach when she discusses *relational* functions. See (Millikan, 2002; Millikan, 1984).

[7] Sometimes, the experience of some activity *is* the end, of course. Roller skates are for those that enjoy skating, for instance, and here there seems to be some overlap between the use plan and the functional goal. But the action "skate" is distinct from the propositional function "is skating". In order to realize the end, that is, in order to be skating, one must perform the action, that is, he must skate.

[8] Those in the Continental tradition may doubt that this distinction is plausible, since capabilities depend on intentional features of the interpreter as purposes do. Nonetheless, we hope that the distinction introduced by Sorabji and discussed below will still serve our purpose. Even if capability claims depend on one's goals, desires, purposes, and so on, not all such claims are *functional ascriptions* in Sorabji's sense.

[9] McLaughlin also provides there an insightful discussion of the difference between the use of explanatory functions in science and artifact (i.e. instrumental) functions.

[10] Perhaps a fallen log is best seen as an example of a token-level function bearer, rather than a token of a function-bearing type. Nonetheless, let us continue with this example as it serves our examination of Sorabji's rule.

[11] Of course, type T may have other functions and thus T-tokens would be *for* these tasks as well.

[12] One may associate a concept of normal token with each function ascription, rather than each type. Thus, a claw hammer with a broken claw is normal with respect to pounding nails, but not with respect to prying nails. We omit this complication for now.

[13] Some televisions use batteries instead of power cords. When he becomes acquainted with televisions of this type, he will consider that there are at least two different types of televisions: battery-powered and outlet-powered. Each of these types have a corresponding concept of normality, as well as use plans and context specifications that differ in minor ways.

[14] This difference is also discussed in (Houkes, 2006).

[15] Paul Davies rejects this claim in [2000]. He argues that etiological theories are no more capable of dealing with malfunction than Cummins's theory.

[16] Neander refines this rough definition later, by specifying that "what it was selected for" should be interpreted as the "lowest level of description" applicable.

[17] (Franssen, 2006) presents an interesting alternative account: " 'x is a malfunctioning K' expresses the normative fact that x has certain features f and that because of these features, a person p has a reason not to use x for K-ing." This provides a broader—perhaps *too* broad—notion of malfunction, but we are here interested primarily in malfunction-as-inability to perform and so we will largely neglect Franssen's notable approach. It is worth noting that, if a token is malfunctioning in the sense we define below, it is also malfunctioning in Franssen's sense, but the converse does not hold.

[18] See (Franssen, 2006) for a distinction between bad design and malfunction.

References

Cummins, R.: 1975, 'Functional Analysis'. *Journal of Philosophy* **72**(20), 741–765.

Davies, P. S.: 2000, 'The Nature of Natural Norms: Why Selected Functions are Systemic Capacity Functions'. *Noûs* **34**(1), 85–107.

Franssen, M.: 2006, 'The Normativity of Artefacts'. *Studies in History and Philosophy of Science* **37**, 42–57.

Houkes, W.: 2006, 'Knowledge of artefact functions'. *Studies in History and Philosophy of Science* **37**, 102–113.

McLaughlin, P.: 2001, *What Functions Explain: Functional Explanation and Self-Reproducing Systems*. Cambridge University Press.

Millikan, R. G.: 1984, *Language, Thought and Other Biological Categories*. The MIT Press.

Millikan, R. G.: 1989, 'In Defense of Proper Functions'. *Philosophy of Science* **56**, 288–302.

Millikan, R. G.: 2002, 'Biofunctions: Two Paradigms'. In: A. Ariew, R. Cummins, and M. Perlman (eds.): *Functions: New Essays in the Philosophy of Psychology and Biology*. Oxford University Press, pp. 113–143.

Neander, K.: 1991, 'The Teleological Notion of Function'. *Australasian Journal of Philosophy* **74**, 261–301.

Neander, K.: 1995, 'Misrepresenting & Malfunctioning'. *Philosophical Studies* **79**(2), 109–141.

Preston, B.: 1998, 'Why is a Wing Like a Spoon? A Pluralist Theory of Function'. *The Journal of Philosophy* **95**(5), 215–254.

Schurz, G.: 2001, 'What is 'Normal'? An Evolution-Theoretic Foundation of Normic Laws and Their Relation to Statistical Normality'. *Philosophy of Science* **68**(4), 476–497.

Sorabji, R.: 1964, 'Function'. *The Philosophical Quarterly* **14**(57), 289–302.

Vermaas, P. E. and W. Houkes: 2003, 'Ascribing Functions to Technical Artefacts: A Challenge to Etiological Accounts of Functions'. *British Journal for Philosophy of Science* **54**, 261–289.

Vermaas, P. E. and W. Houkes: 2006, 'Technical Functions: a Drawbridge Between the Intentional and Structural Natures of Technical Artefacts'. *Studies in the History and Philosophy of Science* **37**, 5–18.

von Wright, G. H.: 1963, 'Practical Inference'. *The Philosophical Review* **72**(2), 159–179.

Wright, L.: 1973, 'Functions'. *Philosophical Review* **82**, 139–168.