Instrumental Artifact Functions and Normativity

Jesse Hughes[*]

February 17, 2011

Abstract

Artifact functions have a decidedly practical character: We are interested in an artifact's function so that we may use it to achieve our ends. The notion of instrumental functions has recently been introduced in order to emphasize the practical aspects of artifactual functions. In this chapter the normative aspects of instrumental functions are investigated, including both agent norms (concerning how and when an artifact should be used) and artifact norms (concerning well- and malfunctioning artifacts). We close with a brief discussion of non-prescribed artifact uses, when an artifact is used other than as intended.

1 Introduction

Artifacts are inherently practical things, intended to be used to achieve certain kinds of ends. This is, after all, what we mean when we speak about artifactual functions—that things of this sort are good *for* something. Or, to

[*]Philosophy Dept., Bentley University, Waltham MA. Email: `jesse@phiwumbda.org`

put the matter differently, that things of this sort are good for *something*—that is, that the function bearer is good at achieving certain related classes of ends. To possess a function is to be suited for certain specified things.

Moreover, there is a natural relationship between artifactual functions and practical reasoning. More often than not, we are interested in knowing an artifact's function so that we will know how it can be used to achieve various ends which we have now or may later adopt. A sailor is interested in the function of his tiller because he needs to know how to use it to steer the boat. Functional knowledge thus has practical consequences: a sailor who knows how to steer using the tiller has a much greater likelihood of reaching his destination than the unfortunate soul who finds himself alone on a boat without the same functional knowledge.

In other terms, the sailor has gained a clearer understanding of the means available to reach his goal than the marooned incompetent. As a result, he is more likely to succeed, but this knowledge also comes with normative consequences: if the sailor fails to use the tiller in the proper manner, then we should fault his judgment. If our marooned incompetent, on the other hand, does not use the tiller properly (and so, does not reach his destination), it is surely not his fault.[1] He does not know how to steer, while the sailor does. To put this obvious point clearly: what constitutes rational behavior in each circumstance depends on what the agent knows about the circumstance, and this includes knowledge about relevant artifact functions.

Nonetheless, the prevailing theories of function are not well-suited for understanding the normative consequences of functional knowledge. Philosophical discussions of functions have largely focused on their contribution to theoretical analyses, not practical reasoning. This is because much of the philosophical interest in functions is due to their recent reintroduction in bi-

[1]In this dire circumstance, we may well think that the marooned incompetent should *try* to learn how to steer on his own, if need be, but that is another matter and his failure to teach himself steering is not comparable to the sailor's failure to steer when he already knows how.

ology. Consequently, functional theories aim either to explain the presence, prevalence or persistence of function-bearing features (especially biological), as in Wright (1973); Millikan (1989); Neander (1991), et al., or to explain a capacity exhibited by a larger system by a component-wise analysis, as in Cummins (1975). In both cases, functions are useful for the explanations they provide, i.e., for their contributions to our *theoretical* knowledge rather than for any more directly practical consequences.

Because we are interested in both agent and artifact norms, rather than in explanations, *instrumental function* (introduced in Hughes (2009a)) is a more natural concept for our analysis. Instrumental functions are intended to simplify the step from functional claims to practical consequences, by focusing precisely on the relevant practical aspects of certain kinds of functional knowledge. The previously mentioned theories of functions serve a different role: they are primarily concerned with the use of functions in scientific contexts. Instrumental functions, on the other hand, emphasize functional claims in practical contexts—interactions between user and artifact, the connection between artifact performance and normative claims (such as whether the artifact is well- or malfunctioning), function creation and publication in engineering, and so on. For a more thorough comparison of Millikan- and Cummins-style functions to instrumental functions, see ibid.

2 Instrumental functions

What are the features of (certain) function ascriptions that allow one to derive clear practical consequences? That is, how does one step from a claim such as, "Bolt cutters are for removing padlocks," to an intention to use *this* pair of bolt cutters in a particular way so that *this* particular padlock can be removed?

We take for granted an instrumentalist model of practical reasoning, in which such reasoning combines an existing desire with beliefs about (relevant) causal relations to yield an action, an intention to act, or a normative judgment regarding an action, depending on the author and context (we tend to express the conclusion as a normative judgment). To take an example from von Wright (1963):

$$\frac{\text{I want to make the hut habitable.} \qquad \text{Unless I heat the hut, it will not become habitable.}}{\text{Therefore, I must heat the hut.}} \tag{1}$$

In this example, the first premise is a claim that I desire to make a particular proposition to be true—that is, that I want to attain a particular *end*. The second proposition expresses a means-end relation, namely that heating the hut is a necessary means to making it habitable. The conclusion here expresses that practical rationality requires that I heat the hut.

What role, then, should functional knowledge play in an instrumentalist account? How can knowledge about an artifact's function be used in a syllogism like the above? By and large, we expect that the effect must be in the second premise—instrumental functions matter to us because they provide new information about means to our ends. Artifacts are valuable to users because they can be used to attain existing ends, not because they generate new ends.[2]

Thus, the first bit of practical information conveyed by a functional claim is that the relevant artifact can be used as a means to its functional end. To continue with our previous example, if Ethel knows that bolt cutters are for

[2]That's not to say that functional knowledge *never* results in new desires. At least some advertising aims to create consumer desires by providing functional knowledge ("It slices! It dices! It bathes the room in the soothing aroma of fresh peaches!"). We want—and are persuaded to want—artifacts because of their functions. Nonetheless, this relationship between artifact functions and desires is tangential to our investigation here, though it may well be central to an account of innovative engineering. See also the discussion of maieutic ends in engineering in Hughes (2009b).

removing padlocks[3], then she may reason thus:

$$\frac{\text{I want to remove this padlock.} \quad \text{Bolt cutters are for removing padlocks.}}{\text{Therefore, I have a reason to use bolt cutters.}} \tag{2}$$

Our second premise is very different than the corresponding premise in von Wright's syllogism, which alleged that heating the hut was a *necessary* means to our agent's end. In this case, on the contrary, all we know that bolt cutters *may* remove a padlock. Bolt cutters are not, presumably, the only way to remove padlocks.

As a consequence, the conclusion here is weaker than that of von Wright's syllogism (1). We do not conclude that Ethel *must* use the bolt cutters, since there may well be other means capable of removing the padlock. She may achieve her goal by one of these other means. Nonetheless, knowing that bolt cutters can be used thus gives Ethel a reason to use bolt cutters.

Let us make the reasoning more explicit. If one knows that

Bolt cutters are for removing padlocks. (**FN**)

then one may conclude that

Using bolt cutters is a sufficient means to removing padlocks. (**ME**)

We take statement (**FN**) to be a(n incomplete) functional ascription and statement (**ME**) a means-end relation[4] which apparently follows from (**FN**).

[3]Of course, bolt cutters have other functions, such as cutting through chain link fences, but let us grant that removing padlocks is *a* function of bolt cutters.

[4]Following von Wright (1963), we adopt the convention that means are *actions*, such as using an artifact, and ends are states or conditions one may desire to attain.

Thus, we may restate (2) in more explicit means-end terms.

$$\frac{\text{I want to remove this padlock.} \quad \text{Using bolt cutters is a sufficient means for removing padlocks.}}{\text{Therefore, I have a reason to use bolt cutters.}} \tag{3}$$

This, indeed, seems a promising reconstruction of the sort of reasoning we use. I infer the means-end relation (**ME**) from the functional ascription (**FN**) and conclude that I therefore have reason to use bolt cutters.

And yet, the conclusion does not appear to be a normative judgment regarding a particular action, but rather a judgment regarding a *kind* of action. We have not concluded that we should use *this* pair of bolt cutters or *that* pair, but rather that we should use *a* pair. When it comes time to act, however, I must proceed with a particular pair, and not an unspecified pair. We need to step from the general norm here expressed to a norm regarding a particular, clearly specified act.

There is a similar, more subtle general/specific mismatch in the premises. The first premise expresses a desire to remove a particular padlock, while the second expresses that our tool is capable of removing padlocks in general. If (**ME**) actually meant that the bolt cutters in question could remove *any* padlock, then clearly the reasoning goes through, since we could replace the second premise with its consequence:

Using bolt cutters is a sufficient means to removing *this* padlock. (**ME′**)

Alas, it is unreasonable to think that (**ME**) really means that *any* pair of bolt cutters is sufficient to remove *any* padlock. One should always use the right tool for the job, and there is good reason that bolt cutters come in different sizes. A well-documented bolt cutter has a specific maximum capacity it can cut through—the higher the capacity, the more expensive, larger and heavier the pair, generally speaking. What Ethel wants to know is that

Using *this* pair of bolt cutters is a sufficient means to removing *this* padlock. (**ME**″)

Insofar as she can determine that it is, then she has reason to use this pair. Thus, in what circumstances does the functional ascription (**FN**) justify the inference to (**ME**″)?

It should be clear at this point that the sentence written as (**FN**)—what we often take to be the functional claim—is incapable of supporting the inference here. That sentence does not include the information needed for our conclusion, and so it seems that the sentence (**FN**) leaves implicit some of the functional knowledge necessary to draw reliable practical conclusions. In fact, in order to conclude (**ME**″)—or even to determine that this pair of bolt cutters is a *relevant* consideration—one needs answers to the following questions regarding the functional claim:

(a) What kind of function-bearers is this about (and is this pair of bolt cutters one of them)?

(b) What result should they produce?

(c) In what situations should they work (and is this situation one of those)?

In addition, in order to actually act appropriately, we must answer one further question:

(d) *How* do I use it?

These questions correspond to the primary features of instrumental functions, respectively:

(A) the artifact type,

(B) the functional goal,

(C) the specification of contexts of use and

(D) the use plan.

We will give a brief overview of each of these in turn. For more details on these features, see (Hughes 2009a).

Artifact type

For our purposes, instrumental functional ascriptions specify the function of function-bearing types—especially, though not exclusively, artifact types. Even in informal functional talk, the artifact type is typically explicit. For instance, it is clear that "Bolt cutters are for removing padlocks" connects a functional goal (removing padlocks) to an artifact type (bolt cutters).[5]

Artifact types may be broad or narrow, depending on the context. We may speak of the function of bolt cutters generally, or the function of heavy-duty, steel-handled, 24" length Brand X bolt cutters. The breadth of the artifact type helps determine the specificity of the remaining instrumental function characteristics. We may say, for instance, that the Brand X bolt cutters are for removing padlocks with a shackle no greater than 5/16" diameter, a much more specific claim than we make regarding bolt cutters generally. The breadth of the artifact type thus acts as a *level of abstraction*, in the sense of (Floridi and Sanders 2004).

There are natural restrictions on the breadth of the artifact type. Because instrumental functions include use plans, functional goals and contexts of use, the artifact type must be narrow enough so that each token is used in the same way, to achieve the same goal and in the same circumstances. There are,

[5]In unusual circumstances, one may find functions that apply to particular tokens rather than a broader type. We are here interested in type-level instrumental functions, however, since they are more relevant for exploring artifact normativity. See (Hughes 2009a).

for instance, many different kinds of wine bottle openers, from a traditional corkscrew to the syringe and air-pump variety to cordless electric openers with push-button interfaces. Clearly, the use plans for these disparate types are very different: you do not use an electric bottle opener in the same way you use a corkscrew. Thus, the type “wine bottle opener” is too broad for our interests here, because its tokens are associated with different use plans. The artifact type must be narrow enough so that its tokens share the same fundamental function characteristics.

The type should also be broad enough so that we do not distinguish tokens of the same basic design. Two electric bottle openers of the same make and model may differ because they were manufactured at different locations or at different times. We should surely treat these as tokens of the same type, regardless of our level of abstraction. On the other hand, a 2009 (model year) Jetta may be considered a different type than a 2010 Jetta, since the two differ in more ways than just manufacture date. We may distinguish the two model years as different types or identify them as a single type (Jetta), depending on our needs and interests.

We use T to denote the artifact type of a function ascription.

Functional goal

Instrumental functions are practical. They express the suitability of using a kind of artifact in order to bring about a certain state of affairs. In terms of means-end reasoning, the functional goal identifies certain ends which are attainable by using a token of the appropriate type in the appropriate situations and in the appropriate manner. This capacity is the source for the instrumental value of an artifact type—we value bolt cutters *because* they remove padlocks (or cut through fences, and so on). If bolt cutters had neither this nor any other instrumental function, then they would have no instrumental value as a type, although particular tokens may still be used (and hence appreciated) in ad hoc manners (as paperweights in remarkably breezy offices, say).

We use φ to denote functional goals, since the state of affairs towards which our usage aims may be expressed as a propositional function. We do not assume, however, that such goals are Boolean, that is, that one either realizes the goal or not. Some goals (such as, "Stopping the car quickly") can be satisfied to greater or lesser degree. We also do not assume that the user can reliably determine whether (or to what degree) the goal has been realized, but we *do* assume that there is a fact of the matter involved. When I take antibiotics, I have no way of knowing whether they produce the desired effect (even if I feel better, that may be a result of my natural immune system rather than the antibiotics), but either they were effective or not (or were effective to some degree), regardless of my personal ignorance.

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: bolt cutters are not for removing particular padlocks, but for removing padlocks (or padlocks of a particular type, say, with suitably small shanks). This functional goal may be crudely represented as $\mathbf{Remove}(x)$, where x is a variable ranging over padlocks of the appropriate type. A particular use of the bolt cutter for this purpose will aim at removing a particular padlock s and so will be successful just in case, afterwords, $\mathbf{Remove}(s)$ is true.[7]

Contexts of use

Artifacts are designed to work in particular circumstances. A particular design of bolt cutter, for instance, is intended to cut through shanks of a particular range of diameters. Complete knowledge of an artifact's function includes not only *how* but *when* the artifact ought to be used. Restrictions

[6] In this respect, our account is similar to Millikan's *relational* functions. See Millikan (2002, 1984).

[7] For a development of instrumental functions in a more formal setting, see (Hughes 2005).

on the intended *contexts of use* place limits on when one can reasonably expect to achieve his goals by using the artifact. A 24" pair of bolt cutters is intended to cut through a 5/16" (or less) padlock shank. When I try to cut through a 1/2" shank and fail, I cannot complain that this pair of bolt cutters has not done what it should.

Thus, each instrumental function comes with a set C of situations in which the artifact type is intended to be used—the *specification of contexts*. A particular usage occurs in a particular context c. 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. When we use a pair of bolt cutters to remove a padlock, the padlock s is part of the context of use and the use is successful if and only if it realizes **Remove**(s).

(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.

Use plan

Finally, instrumental functions involve particular actions. When we say that bolt cutters are for removing padlocks, we mean that there is a particular way to use bolt cutters in order to effect that end. Instrumental functions come with *use plans*.[8] One does not know an instrumental function (in a

[8]See (Houkes 2006; Vermaas and Houkes 2006).

practical sense) unless she has some idea *how* to use the artifact in order to realize its functional goal.

We denote the use plan α. Like the functional goal, a typical use plan includes typed parameters. For instance, using a pair of bolt cutters to remove a padlock involves positioning the cutters so that the shank of the lock is between the blades and pushing the handles together. If x ranges over tokens of type **BoltCutter** and y over **Padlock** (of the appropriate size), then our simple user plan may be represented as $\mathbf{Position}(x, y); \mathbf{Close}(x)$. A particular application of a pair b of bolt cutters to a particular padlock s would be denoted $\mathbf{Position}(b, s); \mathbf{Close}(s)$.

In practice, use plans may be either explicit and detailed or vague and broad. The amount of detail depends, in part, on the current level of abstraction. They may include conditional actions ("If the stapler is empty, load it."). But in each case, they describe what one should *do*. Thus, use plans provide the *means* for our means-end analysis. Roughly, 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.

From functions to means

As we have argued previously,[9] instrumental functions are characterized by the four features described above. Thus, we will use the tuple $\langle \varphi, \alpha, C, T \rangle$ to denote the instrumental function with artifact type T, functional goal φ, context specification C and use plan α. An instrumental function ascription in its essence, then, entails the following:

In situations c satisfying C, one can use a T-token t as prescribed by $\alpha_{c,t}$ in order to realize $\varphi_{c,t}$. (**FM-1**)

At least, this *would* be the case, if not for the fact that different T-tokens may behave differently. Bolt cutters can be sharp or dull, and sharp bolt

[9](Hughes 2009a).

cutters are more likely to achieve their end than dull.

For this reason, we introduced the concept of "normal tokens" in (Hughes 2009a). Given a functional ascription $\langle \varphi, \alpha, C, T \rangle$, a *normal token* t of type T is one with the physical features required to realize φ in the manner intended—that is, one which has the physical features specified by the artifact type's *design*[10] .

With this notion at hand, we amend (**FM-1**) as follows. Given a instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$, one may conclude that

In situations c satisfying C, one can use a *normal* T-token t as prescribed by $\alpha_{c,t}$ in order to realize $\varphi_{c,t}$. (**FM-2**)

In the next section, we will begin to unpack some of the normative consequences of this entailment. Note, however, that (**FM-2**) does *not* entail the instrumental function ascription $\langle \varphi, \alpha, C, T \rangle$. The statement (**FM-2**) is purely about *capabilities* of (normal) T-tokens. Instrumental functions, on the other hand, include an *intentional* aspect. Not only must T-tokens be capable of realizing φ as specified above, but they must be *valued* for this reason. Gasoline-fueled internal combustion engines can be used to contribute greenhouse gases to the atmosphere, but this is not what they are *for*. Loosely speaking, in order for φ to be the *function* of type T, someone must value T-tokens for their capacity to realize φ.

For a detailed discussion of the teleological nature of instrumental func-

[10]Of course, not every artifact type has a formal specification, but we nonetheless presume that there is a sense of normal token available for each type—more precisely, for each *function*. Carpenter's hammers are for pounding nails. They vary widely in their specifications, but in every case, they have a handle roughly perpendicular to a head. The head should be strong enough to strike an object without fracturing or coming loose from the handle. A hammer which fails to satisfy these basic restraints is certainly not a normal carpenter's hammer and is unlikely to serve its function of pounding nails well.

In this and other cases, one may infer what features are relevant for "normalcy", by analyzing how the type is expected to realize its function. A token with the requisite features is normal, and one without is not. As one might expect, designers, engineers and others with deep technical knowledge will have a more precise grasp of what counts as a normal token than everyday users.

tions, see (Hughes 2009a).

3 Prescribed usage: reasons for action

We see now the most direct, practical consequences of instrumental functions. Knowing that an artifact type has an instrumental function entails a family of associated means-end relations. This, in turn, produces reasons for the agent to act appropriately, assuming that he wants to realize the functional goal in some suitable context. As with other (weakly sufficient, in the terminology (Hughes, Kroes, and Zwart 2005)) means, the normative conclusion of such practical reasoning is fairly weak: the agent simply has a reason to use the handy artifact token, but this reason is defeasible. It is not irrational to forgo the usage, assuming that other means exist to achieve his desired end.

We may now give the explicit syllogism for drawing practical consequences from functional knowledge. In order to sketch the reasoning in natural language, we will use "T-tokens are for realizing φ," to express the instrumental function $\langle \varphi, \alpha, C, T \rangle$, leaving C and α implicit.

$$\begin{array}{l} \text{I want to realize } \varphi_{c,t}. \\ T\text{-tokens are for realizing } \varphi. \\ t \text{ is a normal token of type } T. \\ c \text{ is a situation satisfying the specification } C. \\ \hline \text{Therefore, I have reason to do } \alpha_{c,t}. \end{array} \qquad (\textbf{PU})$$

Let us call a syllogism of this form a *prescribed usage* of t. We may illustrate this with our padlock example once again. In this case, I may reason thus:

> I want remove this padlock.
> 24” bolt cutters are for removing small padlocks.
> This artifact is a normal pair of 24” bolt cutters.
> This padlock is a small padlock.
> ———
> Therefore, I have reason to use this artifact to remove this padlock.
>
> **(PU$_b$)**

If I believe each of the premises is true, then I cannot help but accept the conclusion.

Note, however, that I am not practically irrational if I choose *not* to use this pair of bolt cutters, although I *am* irrational if I do nothing (and intend to do nothing in the future) in order to remove the padlock, all the while still desiring its removal.

My knowledge of the function of bolt cutters thus allows me to infer that using bolt cutters are a means to certain ends. Typically, they are one means among many. Depending on my skills and other tools available, for instance, I may consider picking the lock as an alternative means to my end. I may have both 24” and 36” bolt cutters at my disposal and be free to choose between them. A rational agent in this situation has several options. How shall he decide among them?

This decision problem is not, of course, unique to reasoning about artifacts. In most situations in which we pursue a given end, we have choices to make about the means to that end. We do not aim, therefore, to give a full account of how one selects (or ought to select) one means among many here, but instead will illustrate a few of the relevant factors (reliability, effectiveness and costs and side benefits) that play especially important roles in artifact selection.

Reliability. Given an end φ, an action α is a more *reliable* means to that end than another action β if doing α is more likely to realize φ than doing β.

Thus, reliability is a measure of the likelihood of success. In this situation, if α is otherwise equivalent to β, then the agent has a stronger reason to do α than to do β. Suppose, for instance, that either 24" or 36" bolt cutters may remove the padlock at hand, but the longer bolt cutters are more likely to succeed than the shorter, due to mechanical advantage. In this case, if I have both sizes of bolt cutters available and I am otherwise indifferent between the two, then I have a stronger reason to use the 36" bolt cutters rather than the 24", because the former are more reliable.

Effectiveness. In the same situation, if α realizes φ to a greater degree than β, we say that α is more *effective* then β. All other things being equal, this gives one a stronger reason to do α rather than β, provided the difference in degree *matters* to our agent. On my folding bicycle, for example, I can use either a hand brake (which brakes the front tire) or a backpedal brake (which slows the rear). Both brakes have the same functional goal, namely, to slow or stop the bike in a timely manner, but the handbrakes slow the bike more quickly than the backpedal brake. Sometimes, I strongly desire to stop the bike quickly. On those occasions, I use the hand brakes (or, better, both brakes simultaneously). On other occasions, a gentle stop is sufficient and so I have no stronger reason to use the hand brake over the backpedal brake.

Cost and side benefits. The bulk of practical deliberation, however, concerns other features that distinguish α from β. Actions have side effects that make them attractive or unattractive to the agent. Those effects which are unattractive we refer to as *cost*. This includes literal monetary cost, energy usage, pollution or noise produced and so on. Other effects are valued by the agent, and we refer to these as *side benefits*. This includes simple enjoyment in performing the action, the achievement of other ends pursued by the agent and so on. Such considerations often play the determining role in selecting an appropriate action. If I do not need to slow my bike too quickly, I prefer to use the backpedal, because the hand brake can increase the risk of losing control of the bicycle (since it grabs the front wheel of the bike, rather than the rear). Similarly, a competent sailor who needs to arrive in Martha's Vineyard from Cape Cod in a day or so may choose to sail rather than take the (more reliable and faster) ferry, because he enjoys the ride. Cost and side benefits are obviously a crucial consideration in selecting one's course of action.

This is a broad overview of the simplest practical consequences of instrumental function. Knowing the function of an artifact provides one with certain associated means-end relations. These relations are weakly sufficient, in the sense that they assert proper artifact usage *may* (in some cases, *will*) realize the functional goal. The practical normativity involved in such knowledge is fairly weak: it gives the agent a reason to use the artifact to achieve his ends in appropriate circumstances, but the strength of the reason depends on comparing the reliability, effectiveness, cost and side benefits to other alternatives available to realize that same end.

4 Prescribed usage: what the artifact should do

We turn our attention now to norms regarding artifact tokens in a prescribed usage. Given the instrumental function $\langle \varphi, \alpha, C, T \rangle$, a token t is "supposed to" be capable of realizing φ in a prescribed usage. This fundamental normative claim allows one to distinguish between good (i.e., working or well-functioning) and bad (i.e., malfunctioning) T-tokens.

We should note here that this distinction is about *capabilities* and not *actual performance*, although the two are certainly related. A perfectly good token may, in a prescribed usage, fail to bring about φ. Some artifact types are not expected to realize their functional goals in every prescribed usage. An anti-aircraft missile, for instance, may miss its target although it is working properly. Thus, the fact that *this* missile failed to hit its target is not, by itself, proof that this missile is malfunctioning. The question is whether this missile is *capable* of hitting its target—or, more precisely, whether it is as likely to hit its target as it ought to be.

This raises the question: how likely ought success to be? A properly sharpened pencil that fails to leave a mark on writing paper seems to be a bad (not to mention fairly puzzling) pencil, while the failure of an anti-aircraft missile to strike its target is acceptable. Thus, and not surprisingly, our answer must depend on the artifact type.[11] A well-functioning token is as *reliable* as normal tokens of the same type in the same context of use.

Similarly, a well-functioning token is as *effective* as a normal token of the same type in the same context of use. If a car's brakes are incapable of stopping the vehicle as quickly as a "normal" set of brakes would, then they are not functioning as well as they should.

Thus, we propose the following definition of malfunction:[12]

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

A token which is not malfunctioning is *well-functioning*.

Note that this is a fairly narrow definition of malfunction. We treat only the negative aspects of malfunction: a token malfunctions when it cannot do what it is supposed to do. There is another side of malfunction, namely, a token may malfunction when it does something it should not do. A car that emits more pollution than it should is malfunctioning, although it nonetheless fulfills its obvious function: it is a reliable means of personal transportation. Let us leave aside this alternative sense of malfunction here, to be dealt with in later work.[13]

[11] And the function of interest, in the case of types with multiple functions.

[12] First discussed in (Hughes 2005) and presented in greater detail in (Hughes 2009a).

[13] Franssen (2006) gives an account of malfunction that effectively includes both negative

It follows from Definition (**Mal**) that normal tokens are well-functioning. One may suppose that the converse is true as well, that "normal" and "well-functioning" are synonyms, but this is not the case. Consider, for example, a set of rabbit ear antennas for a television set. Sometimes, one of the "ears" breaks off, resulting in poor reception, but a moderately resourceful person can fix this by attaching either a wire hanger, aluminum foil or both. The resulting contraption is surely still a pair of rabbit ears, though it is not a *normal* pair. Nonetheless it may work adequately well. If so, it is a non-normal but well-functioning token of its type.

As we said, actual performance does not, by itself, entail malfunction. The fact that a token fails to realize its functional goal in a particular usage does not mean that the token is malfunctioning.[14] Nonetheless, repeated failure is certainly a good reason to suspect malfunction. Suppose, for instance, that I hit the power switch on my television and nothing happens. It does not turn on as I expect it to. How should I react?

Most likely, before concluding that the set is broken, I would try hitting the power button a couple more times. Perhaps it did not engage like it should have (either because the power button itself is a malfunctioning component or because, sometimes, even simple buttons fail to engage as they should). Let us suppose that, again, nothing happens. The set is not functioning as it should, as far as I can see.

I look at the set. I considered it a normal token previously and it still appears normal to my (uneducated) eyes. Perhaps, then, I was mistaken about the context of use. The obvious question is whether the set is currently

and positive senses: " '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." If a car emits too much pollution, then this is a reason not to use the car, and so it is malfunctioning in Franssen's sense.

[14]Note that if normal tokens are expected to *always* realize the functional goal, which may be the case with some particularly simple artifacts, then a single failure would indeed indicate that the token at hand is not as reliable as normal tokens and hence is malfunctioning.

powered. Is it plugged into a powered wall socket? If so (and if nothing else suggests that the context of use is not appropriate), then I must conclude that this is, indeed, a prescribed usage.

At this point, when I have confirmed repeated failure in prescribed usage, I have reasonable evidence of malfunction. I could still be mistaken in this conclusion—the failures could be due to some statistical fluke,[15] for instance, or I could be wrong about the actual context of use—the wall socket may not deliver the right voltage, perhaps. If, as it turns out, the set is malfunctioning, then it is not, contrary to appearances, a normal token of its type. It may well appear normal to me, but I am wrong. A trained engineer or repairman should be able to discover how the set fails to be normal.

Notice that, once I come to the conclusion that the token t at hand is malfunctioning, any usage of t is *non-prescribed*. The syllogism (**PU**) from p. 14 includes the premise that t is a normal T-token, and this premise is false if t is malfunctioning. Thus, let us turn our attention to non-prescribed usage.

5 Non-prescribed uses

In this section, we will sketch some initial considerations about norms and non-prescribed uses. We are not here prepared to give a full defense of these normative claims, but rather to simply give an initial starting point to the discussion of non-prescribed uses and responsibilities.

Suppose that I am aware of the instrumental function $\langle \varphi, \alpha, C, T \rangle$, but that I am considering a usage which is *not* prescribed, in the sense of Sect. 3. There are various ways in which a usage may not fit the form of syllo-

[15] If a normal token realizes its goal with reliability less than one, then no finite number of failures will prove beyond doubt that the token is malfunctioning, though the longer the sequence, the greater the probability that it is due to malfunction rather than statistical happenstance.

gism (**PU**)[16]:

- The token t at hand is not a normal T-token.
- The context of use c does not satisfy the specification of contexts C.
- My planned use of t is not consistent with the use plan α.

In each case, the difference between prescribed and non-prescribed usage is a matter of degree. The non-normal token may be more or less similar to normal tokens, the context may be more or less similar to the specifications and my planned usage may be more or less similar to the accepted use plan.

In general, one expects that the more similar the actual situation to the prescribed usage, the greater the likelihood that the non-prescribed usage will be successful, but the details matter. If I wish to cut a small pipe and my only cutting tool I have at hand is a wood saw, I may be tempted to use the wood saw. Unfortunately, this is not a very good plan. Wood saws are not well-suited to cut metal and the saw will likely be damaged in the attempt. On the other hand, if I want to cut a small piece of wood and all I have is a hack saw, it is not such a bad idea to use the tool at hand. It should suffice for cutting wood. In order to judge whether a non-prescribed usage is reasonable, one relies on his judgment and experience to a greater extent than in the prescribed case. Instrumental functions are, in some respects, like moral rules of thumb: we rely on our functional knowledge to avoid more tedious and difficult reasoning in order to realize our goals. When functional knowledge does not apply to the situation at hand, practical conclusions are harder to draw.

An agent often *does* have good, practical reasons to use an artifact in a non-prescribed manner. Ethel's bolt cutters may be designed to remove only small padlocks while she needs to remove a somewhat larger padlock. In this situation, a rational agent may well conclude that "it's worth a try." She

[16] We omit one possibility: the end which I am pursuing is not a functional goal for *any* artifact type T.

may have other options available—she may, for instance, consider buying a larger pair of bolt cutters—but nonetheless conclude that at least attempting the non-prescribed usage is preferable, since she may always pursue the alternative course on failure.

Note, however, that the norms regarding the artifact differ considerably between prescribed and non-prescribed usage. Ethel's bolt cutters are not *supposed* to remove *this* (somewhat larger) padlock. This is not what they were designed for. Thus, failure in this case is *not* good evidence of malfunction. We should not blame a token for failure to do what it was not designed to do.

The responsibility for non-prescribed usage falls squarely on the user, then, and not the artifact (or its designers). When an artifact fails in a prescribed usage, we may regard it as the "fault" of the artifact. It ought to have worked. This is even clearer in the case of costly failure: Ethel's bolt cutters shouldn't break when used to remove a small padlock. But, if they break when used to remove a larger padlock, then we are more likely to blame Ethel, rather than the cutters. Ethel has acted unwisely by using the cutters thus.

Non-prescribed usage may violate more serious norms than mere practical self-interest. In some cases, non-prescribed usage puts property or life in danger. Airplanes are supposed to be flown by persons trained in the skill. This is part of their context of use. If I attempt to fly an airplane despite my lack of training, then I am certainly engaged in a non-prescribed usage that is not likely to realize my goal—I'm not likely to get where I'm going in one piece. Moreover, I am putting others at risk—especially if I somehow manage to get airborne.

Even experienced pilots may crash a plane, of course, but the moral harm here is more significant. An experienced pilot is using the plane as it should be used. I am using a plane in a non-prescribed manner. Because of this fact, I assume more responsibility for harm done. When we deviate from use prescriptions, we either increase the risk or lower our capacity to assess the risk, or both, and are thus a natural subject of greater moral responsibility.

Of course, these moral judgments do not alleviate the designer of all responsibility. A well-designed artifact will perform reasonably well in circumstances that are close to prescribed usage, and known dangers regarding non-prescribed usage should be communicated to the user. Notwithstanding these issues, the more abnormal the usage, the more responsibility falls on the user's shoulders.

Clearly, there is more to be said regarding norms in non-prescribed usage. We leave a more thorough examination of this issue for later work.

6 Concluding remarks

Knowledge about artifactual functions is inherently practical. Such knowledge produces practical reasons for using the artifacts as well as norms for judging artifact tokens. These features of functional knowledge are essential to understanding how persons reason about artifact usage and artifact performance and traditional functional theories (which focus on theoretical explanations rather than practical reasoning) do not make this connection explicit. We have shown here how the theory of instrumental functions, on the other hand, does yield clear practical consequences. Viewed in this way, functional knowledge adds to the available means to certain ends.

To be sure, one could reason about each particular token individually and come to essentially the same practical consequences. Even if Ethel doesn't know what bolt cutters are for generally, given time and cleverness, she could discover that this particular pair of bolt cutters is well-suited for removing padlocks. Of course, this sort of ad hoc reasoning is inefficient to the point of absurdity. We pass on functional information because it serves as an effective rule of thumb for discovering new means to our ends. If Ethel knows what bolt cutters are for generally and knows that this is a pair of bolt cutters, then she has some idea how they can be used and a defeasible expectation that they will suffice to achieve their goals (assuming that they appear normal).

But we often choose to use artifact tokens in situations outside of their intended use. Such non-prescribed uses can be rationally justified, although the user implicitly accepts more of the responsibility for failure in such cases. The fact that these bolt cutters are for removing small padlocks places a limit on my expectations regarding their performance. Failure to remove a small padlock is some evidence of malfunction, but failure to remove a larger padlock—something they were not designed to do—need not reflect badly on the bolt cutters.

It is in this last area, the division of responsibility for non-prescribed uses, that is most open for new research. To what degree does the user accept responsibility for catastrophic failure in such situations? In part, it must depend on whether the particular non-prescribed use was foreseeable by the designer and the catastrophic failure preventable (either by design or by education of the user). But here, too, the devil is in the details: the manufacturer of a mass-produced and widely used product has an apparently greater responsibility to predict such (mis-)use than the craftmaker or hobbyist who produces only a few. The division of responsibility for non-prescribed uses between user and designer is obviously a subtle and difficult topic which we are unable to address fully here.

References

Cummins, Robert. 1975. Functional analysis. *Journal of Philosophy* 72(20): 741–765.

Floridi, L., and J. W. Sanders. 2004. Levellism and the method of abstraction. The final draft of this paper is available as IEG Research Report 22.11.04, see `http://www.wolfson.ox.ac.uk/~floridi/pdf/latmoa.pdf`.

Franssen, Maarten. 2006. The normativity of artefacts. *Studies in History and Philosophy of Science* 37:42–57.

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

Hughes, Jesse. 2005. Means-end relations and artifactual functions: a sketch. Presented at the *Norms, Reasoning and Knowledge in Technology* workshop, Boxmeer, the Netherlands, available at `http://phiwumbda.org/~jesse/papers/index.html`.

———. 2009a. An artifact is to use: an introduction to instrumental functions. *Synthese* 168(1):179–199.

———. 2009b. Practical reasoning and engineering. In *Philosophy of technology and engineering sciences*, ed. Dov M Gabbay, Anthonie Meijers, Paul Thagard, and John Woods, 375–402. Elsevier B.V.

Hughes, Jesse, Peter Kroes, and Sjoerd Zwart. 2005. A semantics for means-end relations. Presented at SEP 2005, available at `http://phiwumbda.org/~jesse/papers/index.html`.

Millikan, Ruth Garrett. 1984. *Language, thought and other biological categories*. The MIT Press.

———. 1989. In defense of proper functions. *Philosophy of Science* 56: 288–302.

———. 2002. Biofunctions: Two paradigms. In *Functions: New essays in the philosophy of psychology and biology*, ed. André Ariew, Robert Cummins, and Mark Perlman, 113–143. Oxford University Press.

Neander, Karen. 1991. The teleological notion of function. *Australasian Journal of Philosophy* 74:261–301.

Vermaas, Pieter E., and Wybo 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, Georg Henrik. 1963. Practical inference. *The Philosophical Review* 72(2):159–179.

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