A FORMAL REPRESENTATION OF PURPOSIVENESS: TOWARDS A TELEOLOGICAL LOGIC

Avery Archer

1. Introduction

How are we to formerly represent the claim that the purpose of the human eye is to perceive visual stimuli? This is the question I wish to address in this paper.[1] More generally, I will be attempting to describe a Kripke-style modal semantics that has application to teleological objects. Teleology, broadly construed, is the investigation of purpose or design. Let us say that some object is teleological just in case it has an aim, function or purpose (what I will henceforth refer to as an object's "telos"). For example, we may say that the telos of a hammer is to drive nails, and the telos of the human eye is to perceive visual stimuli. Thus, both hammers and eyes may be described as teleological objects. Alternatively, we may say that an object is teleological just in case it displays design. In the case of artefacts, like hammers, the design is due to human ingenuity. In the case of biological systems, like the human eye, the design is due to evolution by natural selection. In sum, the telos of an object is the purpose for which it is designed. The question with which we began, and which will be the focus of this paper, may therefore be rephrased as follows: how may we formerly represent the claim that a particular object or set of objects have a certain telos? This is the question I will seek to address in this paper.

2. Preliminary Proposals

One proposal for formerly representing the concept of purposiveness may be put as follows: Let 'Eye' refer to the set of human eyes, and let 'See' refer to the set of things that perceive visual stimuli. The claim that the telos of the human eye is to perceive visual stimuli may be formerly represented as follows:

[1] It is important to distinguish between two distinct, albeit closely related, projects: (1) arriving at a criterion for determining if some X is purposive, and (2) formerly representing the fact that X is purposive. In this paper, I will only be concerned with (2); assuming that a certain X is purposive, how are we to go about formerly representing this fact. Thus, the framework articulated in this paper is not intended to provide criterion for determining if some particular object or system is purposive.

(**2.1**) $(\forall x)(Eye(x) \rightarrow See(x))$

According to (2.1), for something to be a member of the set of human eyes is sufficient for that thing to be a member of the set of things that perceive visual stimuli. However, (2.1) clearly fails to capture what we mean when we say that the telos of the human eye is to perceive visual stimuli. Cases of blindness represent a counterexample to (2.1), but they do not represent a counterexample to the claim that the telos of the human eye is to perceive visual stimuli. Thus, the former is not equivalent to the latter. The take home point seems to be that the claim that something has a telos allows for exceptional cases, and therefore cannot be represented by the universal quantifier. Another suggestion, which will also prove to be insufficient, is to replace the universal with an existential quantifier. This yields:

(**2.2**) $(\exists x)(Eye(x) \wedge See(x))$

According to (2.2), there is something that is both a member of the set of human eyes and a member of the set of things that perceive visual stimuli. (2.2) offers a clear advantage over (2.1) since it does not require that all members of the set of human eyes also be members of the set of things that perceive visual stimuli. However, (2.2) fails to capture what we mean when we say that the telos of the human eye is to perceive visual stimuli since we can imagine a situation in which the second conjunct of (2.2) fails to obtain (making the formula false) and in which we would still wish to say that the corresponding teleological claim is true. For example, suppose that a global pandemic, such as a virulent eye-infection, rendered everyone on earth blind. In such a case, (2.2) would be false since there would be nothing that was both a human eye and that also perceived visual stimuli. However, we would still wish to say that the telos of the human eye is to perceive the visual stimuli even if a pandemic disease prevented all human eyes from fulfilling their telos. Another way this point may be put is to say that (2.2) fails to preserve a distinction between some set of objects failing to have a certain telos and a set of objects having a telos that they all fail to fulfil.

I believe that (2.1) and (2.2) both fail because they attempt to represent the claim that the human eye has a certain telos by focusing solely on how things are in the actual world. However, I believe that our concept of what it means for something to have a telos is an essentially modal notion; one that appeals to how things are in worlds other than the actual world. It is this proposal that I will now explore.

3. Accessibility and Teleologically Ideal worlds

In attempting to formerly represent the concept of purposiveness I will be taking as my starting point the accessibility relation introduced by Saul Kripke.[2] Within a Libnizian modal framework, to say that ϕ is necessarily true means that ϕ obtains in all (metaphysically) possible worlds. By contrast, Kripke semantics relativises the notion of necessary truth to a subset of possible worlds; namely, the set of accessible worlds. The upshot is that modal statements (it is necessary that ϕ, it is possible that ϕ) need not take the same truth value in all possible worlds. For example, suppose that Δ is the only world accessible from Γ and that Γ and Δ are both accessible from Δ. Moreover, let us suppose that $\Delta \Vdash \phi$ and that $\Gamma \nVdash \phi$.[3] On the present model, it is necessarily true that ϕ relative to Γ since ϕ obtains in all worlds accessible from Γ. However, it is not necessarily true that ϕ relative to Δ since ϕ does not obtain in all worlds accessible from Δ.[4] As has become standard, I will be defining the relation of accessibility as an (uninterpreted) binary relation $\mathcal{R}(\Gamma,\Delta)$ that holds between possible worlds Γ and Δ just in case Δ is *accessible from* Γ, or Δ is an *alternative world* to Γ.[5] I will use the expression "Δ stands in relation $\mathcal{R}$ to Γ" to describe $\mathcal{R}(\Gamma,\Delta)$, and the expression "$\Gamma$ stands in the relation $\mathcal{R}$ to Δ" to describe $\mathcal{R}(\Delta,\Gamma)$.[6] If we let Γ denote the actual world, we arrive at the following fundamental translational schema for Kripke-style possible world sematics:

(**3.1**) $\Box\phi =_{def} \phi$ is true at every possible world Δ such that $\mathcal{R}(\Gamma,\Delta)$.

(**3.2**) $\Diamond\phi =_{def} \phi$ is true at some possible world Δ such that $\mathcal{R}(\Gamma,\Delta)$.

Fitting and Mendelsohn [1998] observe that "once accessibility among worlds has been highlighted as the preeminent characteristic of modal logics, we can interpret the modal operators far afield from their originally intended meanings."[7] When applied to physics, for example, the accessibility relation is construed in terms of

---

[2] Kripke, S., [1963], "Semantical Analysis of Modal Logic I, Normal Propositional Calculi", *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9, pp. 67-96.

[3] I am here presupposing the standard notation for a Kripke model, understood as a triple $\langle G,R,\Vdash\rangle$, where $\Vdash$ is a relation between worlds of G and modal formulas, such that: (i) $w \Vdash \sim\varphi$ if and only if $w \nVdash \varphi$, (ii) $w \Vdash \varphi \rightarrow \psi$ if and only if $w \nVdash \varphi$ or $w \Vdash \psi$, and (iii) $w \Vdash \Box\varphi$ if and only if $\forall u(R(w,u) \rightarrow u \Vdash \varphi)$.

[4] Significantly, Kripke semantics allows for the possibility that a given world may fail to be accessible from itself (as is the case with Γ but not the case with Δ in our preceding example). As we shall soon see, this feature of Kripke semantics will be crucially important when we attempt to formerly represent the concept of purposiveness.

[5] See: Fitting & Mendelsohn [1998], *First Order Modal Logic*. Boston, MA: Kluwer Academic Publishers, p. 12.

[6] See and Cf. Fitting & Mendelsohn [1998], p. 9.

[7] Fitting & Mendelsohn [1998], p. 10.

nomological accessibility. ϕ is said to be nomologically necessary just in case ϕ is true at any possible world w that is nomologically accessible from the actual world, where w is nomologically accessible from the actual world just in case it obeys the physical laws of the actual world. In short, ϕ is true at all possible worlds that obey the physical laws of the actual world. In deontic logic, the accessibility relation is construed in terms of morally perfect worlds. ϕ is deontically necessary (i.e., obligatory) just in case ϕ obtains in all morally perfect worlds and deontically possible (i.e., permissible) just in case it obtains in some morally perfect world.

An important difference between nomological and deontic necessity is that the class of nomologically accessible worlds includes the actual world (since the actual world is a member of the class of worlds that obeys the physical laws of the actual world), but the class of morally perfect worlds does not include the actual world (since the actual world is not a member of the class of morally perfect worlds). One upshot of this difference is that under a nomological interpretation, $\mathcal{R}$ is Reflexive, while under a deontic interpretation, it is not. (I will have more to say about Reflexivity in §4, where I limn a teleological interpretation of $\mathcal{R}$.) A second upshot of this difference is that if we were to restrict the universe to the class of morally perfect worlds, the actual world would be omitted from our universe. Under a deontic interpretation, the accessibility relation enables us to avoid this unwelcome result by allowing for imperfect moral worlds (a class that includes the actual world) in our universe, while restricting deontic access to those worlds that are morally perfect.

Under a teleological interpretation, the accessibility relation may be seen as restricting us to the set of *teleologically ideal worlds* (henceforth, T-worlds), defined as the set of possible worlds in which every telos found in the actual world is realised. To say that every telos found in some world Γ is realised in some world Δ is to say that Δ is *teleologically accessible* from Γ. This yields the following fundamental translational schema for teleological necessity and possibility:

(**3.3**) $\boxed{T}\phi =_{def} \phi$ is true at all T-worlds.

(**3.4**) $\Diamond_T\phi =_{def} \phi$ is true at some T-world.

Complications aside, I hold that the human eye has the telos of perceiving visual stimuli just in case it is teleologically necessary that being a member of the set of human eyes is sufficient for being a member of the set of things that perceive visual stimuli.[8] The concept of purposiveness seems to fall somewhere between the notion of lawfulness in the physical sciences and obligatoriness in moral theorising. Like the notion of physical laws (and unlike obligatoriness) purposiveness is a descriptive concept; it tells us something about the way the world actually is, and not merely

[8] I qualify this claim in §6.

about how the world ought to be. We may identify the descriptive dimension of purposiveness with the fact that the telos of a purposive system is determined by facts about the actual world. For example, in the case of a biological system, its telos is determined by what that system was selected for *in the actual world*. In the case of human artefacts, the relevant telos is determined by the intentions of the human designer *in the actual world*. Thus, just as we can only tell which worlds are nomologically accessible by inquiring about which physical laws obtain in the actual world, we can only tell which possible worlds are teleologically accessible by inquiring into what a biological system was selected for, or what an artefact was designed for *in the actual world*. One upshot is that the class of T-worlds is restricted to those possible worlds in which every teleological object has the same telos (i.e., is designed for the same purpose) as in the actual world.

However, since purposes often go unfulfilled in the actual world, the actual world is not a member of the class of T-worlds. In this respect, purposiveness is like obligatoriness; both concepts construe the accessibility relation in terms of a set of worlds that excludes the actual world. Another way this point may be put is to say that there is a prescriptive dimension to the concept of purposiveness. Thus, purposiveness displays both descriptive and prescriptive dimensions, and a formal representation of purposiveness, if it is to be satisfactory, must be able to accommodate this fact. As nomological and deontic interpretations of $\mathcal{R}$ illustrates, the accessibility relation (in its uninterpreted form) is conceived of at such a high level of abstraction that it may be applied to both descriptive and prescriptive concepts. Teleological logic, which can be seen as combining both descriptive and prescriptive dimensions, exploits this high level of abstraction to yield a formal representation of purposiveness that can accommodate the complexity of the quotidian concept.

4. Toward Teleological Logic Proper

Thus far, I have tried to provide an intuitive feel for a basic teleological logic (henceforth, BTL). We may now introduce some additional regimentation by specifying the syntax of BTL. Let us assume that we have a simple propositional language, $\mathcal{L}$. The alphabet of $\mathcal{L}$ consists of:

(i) a denumerable set Π of propositional variables p, q, r, p_1, p_2, . . .;

(ii) the primitive logical connectives $\top$ (verum), $\bot$ (falsum), ~ (negation), $\boxed{\mathrm{T}}$ (teleological necessity), $\Diamond_{\mathrm{T}}$ (teleological possibility), $\wedge$ (conjunction), $\vee$ (disjunction), $\rightarrow$ (material implication), and $\leftrightarrow$ (material equivalence); and

(iii) the parentheses ( ).

The well formed formulas (wffs) of $\mathcal{L}$ consists of the smallest set Σ such that:

(a) every propositional variable in Π is in Σ,
(b) $\top$ and $\bot$ are in Σ,
(c) If p is in Σ then so are $\sim p$, $\boxed{T}p$ and $\Diamond_{T}\, p$
(d) If p, q are in Σ, then so are $(p \wedge q)$, $(p \vee q)$, $(p \rightarrow q)$ and $(p \leftrightarrow q)$.

The sentences under (a) and (b) are the *atomic sentences* of $\mathcal{L}$. $\top$ and $\bot$ are 0-place connectives; ~, $\boxed{T}$, $\Diamond_{T}$ are 1-place connectives; and all remaining connectives are 2-place.[9] Finally, I propose the following axiom schemata for BTL:

BTL:	A1.	All tautologous wffs of $\mathcal{L}$	(TAUT)
	A2.	$\boxed{T}(p \rightarrow q) \rightarrow (\boxed{T}p \rightarrow \boxed{T}q)$	($\boxed{T}$-K)
	A3.	$\boxed{T}p \rightarrow \sim\boxed{T}\sim p$	($\boxed{T}$-D)
	R1.	If $\vdash p$ and $\vdash p \rightarrow q$, then $\vdash q$	(MP)
	R2.	If $\vdash p$ then $\vdash \boxed{T}p$	($\boxed{T}$-NEC)

It should be clear to the observant reader that BTL is simply modal system D, with the relevant notation amended to express a teleological interpretation. A1 is standard in all normal modal systems. According to A2, if a material conditional holds in all T-worlds, and its antecedent holds in all T-worlds, then the consequent of the material conditional also holds in all T-worlds. This is the K axiom present in all normal modal logics, also known as the distribution axiom.

A3 follows from conditions imposed on the binary relation $\mathcal{R}$, which restricts access to worlds that are teleologically ideal (i.e., possible worlds in which every telos is realised). Given some Kripke frame $\langle \mathcal{G}, \mathcal{R} \rangle$, A3 tells us that for any world Γ that is a member of $\mathcal{G}$, there is some world Δ in $\mathcal{G}$ such that $\mathcal{R}(\Gamma,\Delta)$. This guarantees that there is always a possible world fitting the conditions of teleological accessibility; thus ensuring that there is always a T-world we may refer to when we need to formerly represent a teleological proposition. In addition to A1-A3, BTL includes Modus Ponens, which is represented by R1. When A1 and R1 are combined, they yield the full inferential power of the propositional calculus. R2 tells us that if p is a theorem, then the claim that p obtains in all T-worlds is also a theorem.

Taking BTL as our starting point, and using our quotidian intuitions about purposiveness as a guide, I believe we may assess which axioms should and should not be included in a plausible teleological logic. For example, we know that if Δ stands in relation $\mathcal{R}$ to Γ, such that $\mathcal{R}(\Gamma,\Delta)$, and some world Ω stands in relation $\mathcal{R}$ to Δ, such

[9] Cf. Åqvist [1984], p. 665.

that $\mathcal{R}(\Delta,\Omega)$, then Ω must itself be a T-world. Since (given the definition of $\mathcal{R}$) all T-worlds stand in relation $\mathcal{R}$ to Γ, it follows that Ω stands in the relation $\mathcal{R}$ to Γ, such that $\mathcal{R}(\Gamma,\Omega)$. This means that, under a teleological interpretation, the relation $\mathcal{R}$ is transitive. This is equivalent to the following axiom:

A4. $\boxed{\mathrm{T}}p \rightarrow \boxed{\mathrm{T}}\boxed{\mathrm{T}}p$ ($\boxed{\mathrm{T}}$-4)

Earlier, it was noted that the actual world is not a member of the set of T-worlds. Given a teleological interpretation of the accessibility relation, it follows that the actual world is not accessible from itself. This entails the denial of *Reflexivity*; the frame condition on $\mathcal{R}$, according to which $\mathcal{R}(\Gamma,\Gamma)$ for every Γ that is a member of $\mathcal{G}$. Thus, under a teleological reading of $\mathcal{R}$, the following axiom turns out to be false:

(*) $\boxed{\mathrm{T}}p \rightarrow p$ ($\boxed{\mathrm{T}}$-M)

Moreover, since some non-T-world Γ (i.e., the actual world) may fail to stand in the relation $\mathcal{R}$ with respect to some given T-world Δ, such that $\mathcal{R}(\Delta,\Gamma)$ is false, even though Δ stands in the relation $\mathcal{R}$ with respect to Γ, such that $\mathcal{R}(\Gamma,\Delta)$ is true, the following axiom also turns out to be false:

(**) $p \rightarrow \boxed{\mathrm{T}}\langle\mathrm{T}\rangle p$ ($\boxed{\mathrm{T}}$-B)

The upshot is that under a teleological interpretation, $\mathcal{R}$ is not *Symmetric*. The denial of (**) follows from the fact that the actual world is not a world in which all the purposes found in a given T-world are realised. Consequently, while all T-worlds stand in the relation $\mathcal{R}$ to the actual world, the actual world does not stand in the relation $\mathcal{R}$ to any T-world. In fact, only another T-world Ω can stand in the relation $\mathcal{R}$ to some other T-world Δ since (intuitively) it is only in some other T-world Ω that every telos found in Δ is realised.

However, this still falls short of the claim that $\mathcal{R}$ is *Euclidean*; the frame condition that if $\mathcal{R}(\Gamma,\Delta)$ and $\mathcal{R}(\Gamma,\Omega)$, then $\mathcal{R}(\Delta,\Omega)$. All that has been asserted thus far is that if Ω stands in the relation $\mathcal{R}$ to some world Δ, then Ω must be a T-world. This is consistent with the possibility that Ω fails to stand in relation $\mathcal{R}$ to Δ; the conditional remains true even if its antecedent is false. In order to arrive at the Euclidean frame condition, we must add the requirement that all T-worlds are teleologically accessible from all other T-worlds. This requirement is satisfied on the present account of teleological accessibility. Recall that the telos of an object in a given T-world is determined by the purpose that object was designed to fulfil (either by natural selection or human ingenuity) in the actual world. This means that the fact that a hammer is designed to drive nails in the actual world is both necessary and sufficient for hammers to have the telos of driving nails in a given T-world. Since the actual world determines the telos of objects in all T-worlds, it follows that the telos of objects is the same in all T-worlds. Add to this the condition that the set of teleological objects found in any T-world is identical to that found in the actual world, so that hammers exist in some T-world if and only if hammers exist in the actual world. It follows that the set of teleological objects is the same across all T-worlds. When these three conditions are combined—(i) the condition that the set of teleological objects is the same across T-worlds, (ii) the condition that teleological objects have the same telos across T-worlds, and (iii) the condition that a T-world is one in which every telos is realised—the upshot is that every T-world stands in relation $\mathcal{R}$ to every other T-world. When this observation is combined with the fact that all T-worlds stand in relation $\mathcal{R}$ to the actual world, we arrive at the Euclidean frame condition on $\mathcal{R}$.[10] This yields the following axiom:

A5. $\Diamond_{T} p \rightarrow \boxed{T} \Diamond_{T} p$ ($\boxed{T}$-5)

Moreover, if all T-worlds stand in the relation $\mathcal{R}$ to each other, then all T-worlds stand in the relation $\mathcal{R}$ to themselves. Consider: if T-worlds are possible worlds in which every telos is realised, then every telos found in a given T-world must be realised in that T-world. It follows that for any given T-world, it stands in the relation $\mathcal{R}$ to itself. However, as was noted earlier, the actual world is not a T-world, so that the actual world fails to stand in the relation $\mathcal{R}$ to itself. This, we noted, entails the denial of Reflexivity. However, any world which occupies the second position in the two-place relation $\mathcal{R}(\Gamma,\Delta)$ must (by definition) be a T-world, which means that it must stand in the $\mathcal{R}$ relation to itself. It follows that $\mathcal{R}$ is *Shift Reflexive*, such that if $\mathcal{R}(\Gamma,\Delta)$ then $\mathcal{R}(\Delta,\Delta)$. This yields the following axiom:

A6. $\boxed{T}(\boxed{T}\, p \rightarrow p)$ ($\boxed{T}$-$\boxed{T}$M)

Earlier, we noted that since the actual world does not stand in the relation $\mathcal{R}$ to any T-world (even though all T-worlds stand in the relation $\mathcal{R}$ to the actual world), $\mathcal{R}$ is not Symmetric. Even so, if $\mathcal{R}(\Gamma,\Delta)$ holds for some world Δ, and Ω is accessible from Δ, such that $\mathcal{R}(\Delta,\Omega)$, then Ω must be a T-world. But if Ω is a T-world, and given that all T-worlds are accessible from each other, then Δ must stand in relation $\mathcal{R}$ to Ω, such that $\mathcal{R}(\Omega,\Delta)$. This means that $\mathcal{R}$ is *Shift Symmetric*, such that if $\mathcal{R}(\Gamma,\Delta)$ holds for some world Δ, then $\mathcal{R}(\Delta,\Omega)$ only if $\mathcal{R}(\Omega,\Delta)$. Thus, we arrive at the following axiom:

A7. $\boxed{T}(\Diamond_{T} \boxed{T} p \rightarrow p)$ ($\boxed{T}$-$\boxed{T}$B)

[10] Cf. Fitting & Mendelsohn [1998], p. 10.

Euclidean modal systems are usually assumed to be Transitive, Reflexive and Symmetric, as with system S5. However, while $\mathcal{R}$ is Transitive (given a teleological interpretation), it is not Reflexive and Symmetric. Instead, $\mathcal{R}$ is Shift Reflexive and Shift Symmetric. When Transitivity, Shift Reflexivity, Shift Symmetry and the Euclidean frame condition are added to BTL, we arrive at what may be referred to as Sophisticated Teleological Logic (henceforth STL):

STL:	A1.	All tautologous wffs of $\mathcal{L}$	(TAUT)
	A2.	$\boxed{T}(p \rightarrow q) \rightarrow (\boxed{T}p \rightarrow \boxed{T}q)$	($\boxed{T}$-K)
	A3.	$\boxed{T}p \rightarrow {\sim}\boxed{T}{\sim}p$	($\boxed{T}$-D)
	A4.	$\boxed{T}p \rightarrow \boxed{T}\,\boxed{T}p$	($\boxed{T}$-4)
	A5.	$\Diamond_T\, p \rightarrow \boxed{T}\,\Diamond_T\, p$	($\boxed{T}$-5)
	A6.	$\boxed{T}(\boxed{T}\, p \rightarrow p)$	($\boxed{T}$-$\boxed{T}$M)
	A7.	$\boxed{T}(\Diamond_T\,\boxed{T}p \rightarrow p)$	($\boxed{T}$-$\boxed{T}$B)
	R1.	If $\vdash p$ and $\vdash p \rightarrow q$, then $\vdash q$	(MP)
	R2.	If $\vdash p$ then $\vdash \boxed{T}p$	($\boxed{T}$-NEC)

STL will provide the foundation for the teleological logic limned in this paper. To this axiomatic framework we will add the machinery of *possibilist quantification* (i.e., the semantics for quantifiers with *constant domain* models), and multimodality.[11] In the next section I will attempt to motivate a multimodal teleological logic by considering a number of objections to the version of T-world semantics described thus far.

5. Toward a Multimodal Teleological Logic

One shortcoming of the preceding account of a T-world is that it gives rise to circularity when offered as an explication of the concept of purposiveness. Recall, a T-world has been defined as a world in which every telos (i.e., purpose) is realised. Thus, the concept of a T-world already appeals to the concept of a purpose, and could not feature in a non-circular explication of what it *means* for an object to have a purpose. We may avoid this shortcoming by characterising T-worlds in a way that does not already presuppose the concept of a purpose. To this end, I posit the following definition of a T-world:

[11] A discussion of the difference between varying and constant domain models falls outside the scope of this paper (See: Fitting & Mendelsohn [1998], pp. 95-115 for an in-depth treatment). In the discussion to follow I will be presupposing a constant domain framework.

DEFINITION 5.1: *T-Worlds*
Some world, w, is a T-world just in case for all $i \in w$, if i is designed to bring about some state of affairs n, then i is a sufficient condition for n.

According to DEFINITION 5.1, a T-world is one in which the existence of a teleological object is sufficient to bring about the state of affairs it was designed to bring about. However, we are still left with the task of specifying what it means for an object to be designed to bring about some state of affairs n. I propose separate accounts for biological systems, whose design is due to evolution by natural selection, and artefacts, whose design is tied to the intentions of its human creator.

DEFINITION 5.2: *Design in Organisms*
If i is some feature of organism *O*, then i is designed to bring about n just in case the fact that n increased the inclusive fitness of the ancestors of *O*, and the fact that i facilitates n, explains the existence of i with respect to *O*.

DEFINITION 5.3: *Design in Artefacts*
If i is an artefact, then i is designed to bring about n just in case the creator of i made it with the intention of bringing about n.

Insofar as we take T-world semantics to implicate something along the lines of DEFINITIONS 5.1, 5.2 and 5.3, it provides us with a non-circular way of representing the concept of purposiveness.

However, it should also be noted that the task of explicating the meaning of a concept is distinct from the task of formally representing that concept. Moreover, there is no difficulty in presupposing a concept in an attempt to formerly represent that concept. For example, the attempt to formerly represent lawfulness in physics and obligatoriness in nomological and deontic logic, respectively, both make explicit appeals to the notions of lawfulness and obligatoriness. Likewise, the present attempt to offer a teleological interpretation of the accessibility relation is only concerned with representing the concept of purposiveness rather than explicating the meaning of the concept. What DEFINITIONS 5.1, 5.2 and 5.3 show is that by filling in the details of the teleological interpretation of $\mathcal{R}$, we may arrive at a non-circular explication of the concept. Still, it is important to keep in mind that the two projects—that of explicating the concept of purposiveness and that of formerly representing the concept of purposivess—are orthogonal to each other. Since I am presently only concerned with latter, there is no danger attached to an explicit appeal to the concept of purposiveness when describing the teleological interpretation of $\mathcal{R}$.

Another putative difficulty with characterising T-worlds as possible worlds in which every telos is realised is that it seems to preclude compensatory and/or conflicting purposes. The notion of a compensatory purpose applies to teleological objects that have the telos of "filling in" for when some other teleological object fails to realise its telos. For example, we can imagine a system equipped with an emergency self-destruct sequence that is only initiated if there is a failure in all other safety protocols. If we conceive of T-worlds as worlds in which every telos is realised, then such a self-destruct mechanism will never have the opportunity to realise its telos since there will never be the required failure elsewhere in the system. This suggests that in a given T-world, compensatory purposes are never realised. However, if compensatory purposes remain unrealised, then a T-world cannot really be a world in which every telos is realised. Thus, the possibility of compensatory purposes appears to threaten the concept of a T-world with incoherence.

One possible reply to the above objection, which will ultimately prove inadequate, is to distinguish between cases in which some X does not have the opportunity to fulfil the function for which it is designed and cases in which X is presented with such an opportunity, but fails to do so. Moreover, we may say that a teleological object only fails to realise its telos in the latter case. This suggestion appears to have some intuitive traction since few would regard a safety mechanism as somehow defective simply because it never had the opportunity to perform its function. On this view, some world Γ counts as a non-T-world just in case some X in Γ is presented with an opportunity to realise its telos and yet fails to do so. Thus, far from implying the failure of compensatory purposes, T-worlds virtually guarantee that such purposes are realised by removing the antecedent conditions for their failure. Consequently, the concept of a T-world remains coherent, even if there are teleological objects with compensatory purposes.

The above reply seems right, as far as it goes. But there are at least two reasons for thinking that it does not go very far. First, while it appears to avoid the problem posed by compensatory purposes in the case of artefacts, it is not clear that the reply generalises to biological systems. We can certainly imagine a biological system or process that has the function of "filling in" for some other biological system or process, should the latter fail to realise its telos. Moreover, we can imagine such a system evolving in the actual world since individuals within a biological population with such a system acting as "back-up" would display greater reproductive success than their conspecifics that lacked such "safety nets". However, it is not clear that such a "back-up" system could ever evolve in a T-world. Since the original system will always realise its telos, the presence of a "back-up" system will never offer any evolutionary advantage to its possessor, and will therefore never become subject to selective pressures. The upshot is that there could never be a biological system with a compensatory purpose in a T-world.[12]

Second, the definition of T-worlds as possible worlds in which every telos is realised seems to be at odds with the fact that there may be conflicting purposes. The notion of a conflicting purpose applies to teleological objects whose telos involves preventing some other teleological object from realising its telos. For example, we can imagine a type of antibiotic whose purpose it is to prevent the DNA found in a particular bacteria from performing its replicatory function. Insofar as we define a T-world as one in which every telos is realised, then both the bacterial DNA and the antibiotic cannot coexist in the same T-world. But since we have assumed a constant domain semantics, and since there are possible worlds in which both the bacterial DNA and the antibiotic exist (e.g., the actual world), this appears to throw the notion of a T-world into jeopardy.

I believe we may overcome the challenge posed by both compensatory and conflicting purposes via a multimodal teleological logic, in which the accessibility relation $\mathcal{R}_i$ is indexed to a particular biological system or human artefact, i. Instead of a single "common" accessibility relation $\mathcal{R}$, there is a series $\mathcal{R}_i$, $\mathcal{R}_j$, $\mathcal{R}_k$ . . . , indexed to sets of teleological objects (e.g., the set of human eyes, the set of human ears, the set of hammers). The Kripke frame for the corresponding language, $\mathcal{L}$, in which $\{ \boxed{\mathrm{T}}i \mid i \in \mathrm{I}\}$ represents the set of necessity operators of $\mathcal{L}$, consists of a non-empty set of possible worlds $\mathcal{G}$, and the binary relation $\mathcal{R}_i$, for each $i \in \mathrm{I}$. The satisfaction relation for $\boxed{\mathrm{T}}i$ is defined as follows:

(**5.1**) $w \Vdash \boxed{\mathrm{T}}i\, \phi$ if and only if $\forall u(\mathcal{R}_i(w,u) \rightarrow u \Vdash \phi)$.

According to (5.1), ϕ is true in all T-worlds relative to some biological system or artefact i just in case ϕ holds in any world that stands in the relation $\mathcal{R}$ to some other world. Instead of speaking of T-worlds in which every telos is realised, we now speak of T-worlds indexed to some biological system or artefact, i, such that i always realises the telos for which it was selected or designed in the actual world. We may revise our definition of a T-world to reflect this multimodality:

DEFINITION 5.4: *T-Worlds*

Some world, w, is a T-world indexed to i just in case $i \in w$ and for any m, such that m is a member of i, if m is designed to bring about some state of affairs n, then m is a sufficient condition for n.

[12] The present objection is potentially more devastating than the first, for it not only casts doubt on some compensatory system ***i*** being able to realise its telos in a T-world, but it also casts doubt on the possibility that ***i*** may have a telos at all.

DEFINTION 5.4 allows for some world w to be a T-world relative to some i even if some $j \in w$ fails to fulfil its purpose, so long as $j \neq i$. On this view, the possible worlds in which the antibiotic performs its function constitutes the set of T-worlds indexed to that antibiotic, while the possible worlds in which the DNA of a particular bacterium realises its telos represents the set of T-worlds indexed to that bacterial DNA. Since each teleological object is now indexed to its own set of T-worlds, there can be no conflicting purposes within T-worlds. Thus, the coherence of the concept of a T-world is preserved. The multimodal account also avoids the problem posed by biological systems with compensatory purposes. Since the set of T-worlds indexed to some biological system i includes worlds in which some other biological system j fails to realise its telos, this allows i to increase the reproductive fitness of its host when i serves as a "back-up" system in the eventuality of j failing to realise its telos. The upshot is that on the multimodal account, there is no difficulty posed by cases of compensatory or conflicting purposes. The axioms for the proposed multimodal logic remain the same as those of STL, except for the addition of the indexical, i, which stands for a set of teleological objects.

STL*i*:	A1.	All tautologous wffs of $\mathcal{L}$	(TAUT)
	A2.	$\boxed{T}i(p \to q) \to (\boxed{T}ip \to \boxed{T}iq)$	($\boxed{T}i$-K)
	A3.	$\boxed{T}ip \to {\sim}\boxed{T}i{\sim}p$	($\boxed{T}i$-D)
	A4.	$\boxed{T}ip \to \boxed{T}i\boxed{T}ip$	($\boxed{T}i$-4)
	A5.	$\langle T \rangle ip \to \boxed{T}i\langle T \rangle ip$	($\boxed{T}i$-5)
	A6.	$\boxed{T}i(\boxed{T}ip \to p)$	($\boxed{T}i$-$\boxed{T}i$M)
	A7.	$\boxed{T}i(\langle T \rangle i\boxed{T}ip \to p)$	($\boxed{T}i$-$\boxed{T}i$B)
	R1.	If $\vdash p$ and $\vdash p \to q$, then $\vdash q$	(MP)
	R2.	If $\vdash p$ then $\vdash \boxed{T}ip$	($\boxed{T}i$-NEC)

STL*i* is a multimodal variant of STL, and (when combined with universal and existential quantification) represents the teleological logic I wish to endorse.[13]

[13] The multimodal teleological logic presently on offer is meant to be understood atemporally—to wit, some world *w* is teleologically ideal with respect to some item ***i*** just in case ***i*** fulfils the purpose for which it is designed at all times; past, present and future. If we assume, quite plausibly, that no teleological object ***i*** fulfils the purpose for which it was designed at all times in the actual world, then the actual world remains a non-T-world. However, once the multimodal apparatus is introduced, we may choose to index T-worlds not only to particular biological systems or artefacts, but also to particular times. On such a view, some world *w* is teleologically ideal with respect to some item ***i*** at some time t just in case ***i*** fulfils the purpose for which it is designed, at t. On such a multimodal account, the actual world may sometimes be included among the set of T-worlds, so long as ***i*** fulfils the function for which it is designed in the actual world, at t. This would require that we adjust the axioms of STL, since the accessibility relation, so indexed, would exhibit both Reflexivity and Symmetry. However, at some other time t—a time in which ***i*** fails to fulfil its function in the actual world—the actual world would be excluded from the set of T-worlds. Thus, when we allow for temporal indexing, we allow for varying frame conditions, some

6. Teleological Necessity, Possibility and Impossibility

According to the view presently on offer, we may represent the claim that the telos of the human eye is to perceive visual stimuli by saying that it is teleologically necessary that being a member of the set of human eyes is sufficient for being a member of the set of things that perceive visual stimuli. Expressed in terms of a multimodal T-world semantics, the human eye has the telos of perceiving visual stimuli just in case it is true in all T-worlds indexed to the set of human eyes that being a human eye is sufficient for perceiving visual stimuli. Formerly, this may be represented as follows:

(**6.1**) $\boxed{T}_{Eye}(Eye(x) \rightarrow See(x))$

(6.1) avoids the difficulty that faced (2.1) and (2.2) since the accessibility relation associated with the modal operator restricts the sufficiency claim to T-worlds (indexed to the set of human eyes). This means that the human eye would continue to have the telos of perceiving visual stimuli even if a pandemic rendered all human beings in the actual world blind.

Significantly, none of the variables in (6.1) are bound, which (in the terminology of classical logic) means that (6.1) is a well-formed formula but not a sentence. We can transform (6.1) into a sentence by introducing a quantifier that ranges over all the free variables in (6.1). The particulars of the case will dictate which quantifier we use. In the case of the human eye, the function of perceiving visual stimuli is assumed to be universal.[14] If this assumption is right, then it follows that in any T-world indexed to the set of human eyes, all human eyes would perceive visual stimuli. Thus, we may revise (6.1) to read:

(**6.2**) $\boxed{T}_{Eye}(\forall x)(Eye(x) \rightarrow See(x))$

However, it is not always the case that all members of a certain class can be said to share a certain telos. For example, consider the set of bird's wings. While some bird's wings have the function of facilitating flight, others (like that of the ostritch, emu and penguin) do not. Moreover, the exceptional cases just listed are not flightless due to disease or some anatomical defect, but because of their evolutionary history. Thus, even in T-worlds indexed to the set of all bird's wings, it is not true that all bird's

---

allowing for Reflexivity and Symmetry, and others not. In order to avoid such complexities, I propose that we restrict ourselves to an atemporal multimodal account. Hence my endorsement of STL*i*.

[14] Some caution is called for here. The present claim is not that all human eyes perceive visual stimuli, but that all human eyes have the *function* of perceiving visual stimuli. Cases of blindness represent an instance in which the human eye fails to perform its function, rather than an instance in which the human eye fails to have a function.

wings facilitate flight. We may formerly represent this fact, letting 'Wing' stand for the set of bird's wings and letting 'Flight' stand for the set of things that facilitate flight, as follows:

(**6.3**) $\boxed{T}_{Wing}(\exists x)(Wing(x) \wedge Flight(x))$

The goal of the indexing machinery is to restrict attention to worlds that are teleologically ideal with respect to a single item at a time. This may be described as a local rather than global notion of teleological ideality. However, it does not restrict attention to one telos at a time. The claim that a particular item has more than one telos may be represented as follows: Assuming that i indexes T-worlds to the set of human eyes; that 'Eye' stands for the set of human eyes; that 'See' stands for the set of things that perceive visual stimuli; and 'Tears' stands for the set of things the produce tears; we may represent the claim that the human eye has the telos of facilitating vision and producing tears as follows:

(**6.4**) $\boxed{T}_{Eye}(\forall x)(Eye(x) \rightarrow (See(x) \wedge Tears(x))$

This may be literally read: in all T-worlds (indexed to human eyes), being a member of that set is sufficient for facilitating vision and producing tears. Similarly, if we wanted to say that *some* bird's wing had the telos of facilitating flight and protecting young birds from harm, this may be represented as follows:

(**6.5**) $\boxed{T}_{Wing}(\exists x)(Wing(x) \wedge Flight(x) \wedge Protect(x))$

As with standard modal logic, the modal operator ranges over possible worlds, while the universal and existential quantifiers range over objects within a given world. In the context of T-world semantics, this means that the universal and existential quantifiers only range over objects in teleologically ideal worlds. This is an important result, and one that seems to accord with common sense. For example, suppose that a global avian disease rendered all birds flightless, so that sparrows and hawks were forced to locomote bipedally like ostriches and emus. Still, we may wish to distinguish between those birds that lost the ability to fly due to disease and those that are naturally flightless by noting that wings have the telos of facilitating flight in the former but not the latter. Thus, it may be true that some bird's wings have the telos of facilitating flight even if, in the actual world, no bird's wings actually did facilitate flight. This suggests that the truth of the existential statement depends on how things are in teleologically ideal worlds and not on how things are in the actual world.

The possibility operator in teleological logic is analogous to the claim that something is permissible in deontic logic. To say that X obtains in *some* morally perfect worlds is not to say that X is obligatory in some morally perfect worlds. Rather, it is to say that X is permissible (as opposed to being obligatory). On a deontic reading, X is obligatory only if X obtains in all morally perfect worlds. In other words, obligatoriness is only expressed by the necessity operator. Similarly, to say that in the set of T-worlds indexed to human ears, there are some T-worlds in which being a member of that set is sufficient for being a member of the set of things that facilitate the wearing of glasses does not mean that some human ears have the purpose of facilitating the wearing of glasses. (Purpose is only expressed by the necessity operator.) Rather, it means that the wearing of glasses is consistent or compatible with the human ear fulfilling its purpose (whatever that purpose happens to be). In this respect, teleological possibility is very much like deontic permissibility; it allows us to represent the claim that a certain state of affairs is consistent with an object's telos.

The teleologically possible refers to states of affairs that obtain in some but not all T-worlds. For example, there are some T-worlds in which being a member of the set of human ears is sufficient to facilitate the wearing of glasses—to wit, we can imagine worlds in which the human ear does what it was designed to do (in the actual world) and in which it also facilitates the wearing of glasses. Included among such T-worlds are ones in which human ears have a shape and position similar to that of human ears in the actual world. However, we can imagine possible worlds in which evolutionary history and selective pressures are such that human ears have a shape and position that does not facilitate the wearing of glasses and in which they nevertheless do what they were designed to do (in the actual world). Insofar as human ears successfully accomplish in some world *w*, what they designed to accomplish in the actual world, *w* qualifies as a T-world with respect to human ears. Thus, there are some T-worlds (indexed to the set of human ears) in which the human ear facilitates the wearing of glasses and other T-worlds (indexed to the set of human ears) in which they do not. The upshot is that while facilitating the wearing of sunglasses is compatible with the human ear fulfilling its function, it is *not* the telos of the human ear. We may formerly represent this idea, letting 'Ear' stand for the set of human ears and 'Glasses' for the set of things that facilitate the wearing of glasses, as follows:

(**6.6**) $\Diamond_{T}{}_{\text{Ear}}(\text{Ear}(x) \rightarrow \text{Glasses}(x))$

We may further elaborate on (6.4) with the addition of a quantifier. As before, what quantifier (if any) we choose to employ will depend on the particulars of the case. For example, we can easily imagine a T-world in which all ears facilitate the wearing of glasses. We can formerly represent this possibility as follows:

(**6.7**) $\Diamond_{T}{}_{\text{Ear}}(\forall x)(\text{Ear}(x) \rightarrow \text{Glasses}(x))$

According to (6.5), there is some T-world (indexed to the set of human ears) in which all ears facilitate the wearing of glasses. We can also imagine T-worlds in which only some human ears facilitate the wearing of glasses, in which case the existential quantifier would be apt:

(**6.8**) $\Diamond_{T\,\mathrm{Ear}}(\exists x)(\mathrm{Ear}(x) \rightarrow \mathrm{Glasses}(x))$

Finally, the *teleologically impossible* refers to states of affairs that fail to obtain across all T-worlds. These include all states of affairs that are inconsistent with the successful functioning of a biological system or artefact that falls within the class of objects to which the T-world in question is indexed. For example, in T-worlds indexed to the set of human ears, it is teleologically impossible for a viral ear infection to prevent the human ear from perceiving auditory stimuli. However, in T-worlds indexed to the aforementioned virus there may be worlds in which having human ears is insufficient for perceiving auditory stimuli. Thus, states of affairs that are teleologically impossible relative to one *ί*-class of T-worlds may be teleologically possible relative to another *ί*-class of T-worlds.

7. Conclusion

I wish to conclude by considering a potential objection to the account of T-worlds articulated above. In §5 I introduced a multimodal framework in response to the challenges posed by compensatory and conflicting purposes. This framework requires that we specify the *ί*-class of T-worlds to which a particular teleological object is indexed before we can formally represent the claim that it has a certain telos. But it is not immediately clear that this requirement can be met in cases in which only some members of the *ί*-class have a certain telos. For example, consider the set of bird's wings, which includes some members that have the telos of facilitating flight and other members that do not. In §6 I suggested that if we let 'Wing' stand for the set of bird's wings and 'Flight' stand for the set of things that facilitate flight, the claim that some bird's wings have the telos of facilitating flight may be formerly represented as follows:

(**6.3**) $\boxed{T}_{\mathrm{Wing}}(\exists x)(\mathrm{Wing}(x) \wedge \mathrm{Flight}(x))$

The immediately preceding formalisation requires that we specify the *ί*-class of T-worlds to which the set of bird's wings is indexed. Since T-worlds are defined as worlds in which every telos is realised, the relevant *ί*-class can only include worlds in which every member of the set of bird's wings realises its telos. But since only some bird's wings have the telos of facilitating flight, we seem unable to specify an *ί*-class that includes all bird's wings. The upshot is that we seem unable to represent the claim that some members of a particular set have a telos since the i-class corresponding with any given set can only include those members that not only have a telos, but whose telos is also realised, or at least so the objection goes.

Worlds in which only some members of a particular set have a certain telos are of broadly three kinds. The first kind of world is one in which some members of the relevant set have one telos and other members have another telos. For example, consider a world in which the set of bird's wings only has two members: a pair of eagle wings and a pair of penguin wings. Moreover, suppose that while the pair of eagle wings has the telos of facilitating flight, the pair of penguin wings has the telos of facilitating swimming. In such a world, all the members of the set of bird's wings may be said to have a telos, albeit not the same one. The second kind of world is one in which some members of a particular set have a telos while other members of that same set have no telos at all. For example, consider a world in which the set of bird's wings includes only a pair of eagle wings and a pair of kiwi wings. Moreover, let us suppose that while the pair of eagle wings has the telos of facilitating flight, the pair of kiwi wings is purely vestigial; it has no telos whatsoever. In such a world, it is not only true that only some bird's wings have the telos of facilitating flight, but it is also true that some bird's wings have no telos at all. Finally, there are mixed worlds in which some members of the relevant set have one telos, other members of that same set have a different telos, and still other members of that set have no telos at all. An example of such a world is one in which the set of bird's wings included a pair of eagle wings, a pair of penguin wings and a pair of kiwi wings. Each of the aforementioned worlds represent a different kind of case in which we would wish to say that only some bird's wings have the telos of facilitating flight.

I believe that DEFINITION 5.4 provides us with the theoretical resources we need to accommodate all three of the aforementioned cases. To recap:

DEFINITION 5.4: *T-Worlds*
Some world, w, is a T-world indexed to i just in case $i \in w$ and for any m, such that m is a member of i, if m is designed to bring about some state of affairs n, then m is a sufficient condition for n.

Let us begin with the first kind of case; instances in which some members of a set have one telos and other members of that same set have a different telos. For simplicity, let us suppose that the set of bird's wings only has two members: a pair of eagle wings and a pair of penguin wings. Moreover, let us assume that the pair of eagle wings is designed (by natural selection) to facilitate flight and that the pair of penguin wings is designed (by natural selection) to facilitate swimming.[15] According to DEFINITION 5.4, a T-world indexed to the set of bird's wings is one in which both the pair of eagle wings and the pair of penguin wings realise their respective telos. However, DEFINITION 5.4 does not require that the pair of eagle wings and the pair of penguin wings have the same telos. In other words, a T-world indexed to the set of bird's wings counts as such just in case every member of that set realises the telos for which it was designed, whatever that telos may be. Thus, our definition of a T-world is able to easily accommodate the first case.

Let us now turn to the second kind of case; instances in which some members of a set have a telos and other members of that set have no telos at all. For simplicity, let us again assume that the set of bird's wings only has two members: a pair of eagle wings and a pair of kiwi wings. According to DEFINITION 5.4, a world containing (exhaustively) a pair of eagle wings and a pair of kiwi wings counts as a T-world indexed to the set of bird's wings just in case the following material conditional holds: if the pair of eagle wings was designed to bring about some state of affairs, then that state of affairs obtains, and if the pair of kiwi wings was designed to bring about some state of affairs, then that state of affairs obtains. Since, *ex hypothesi*, the pair of kiwi wings was not designed to bring about any state of affairs, the antecedent of the conditional is not satisfied, making the conditional as a whole true. The upshot is that in a world in which the set of bird's wings only contains a pair of eagle wings and a pair of kiwi wings, its status as a T-world is determined solely by whether or not the pair of eagle wings realises its telos in that world. If it does, then that world counts as a T-world indexed to the set of bird's wings.

Finally, we may imagine a world in which the set of bird's wings includes a pair of eagle wings, a pair of penguin wings and a pair of kiwi wings. This world combines elements from both of the preceding worlds. Such a world would count as a T-world (indexed to the set of bird's wings) just in case the pair of eagle wings and the pair of penguin wings realise their respective telos. Since the pair of kiwi wings have no telos, there is no danger of its telos failing to be realised. This allows us to further clarify what it means for a possible world to be a T-world indexed to a particular teleological object. A T-world indexed to some set of objects is not one in which every member of that set has a certain telos. Rather, a T-world indexed to a set of objects is one in which every member of that set successfully brings about the state of affairs it was designed to bring about, but only insofar as it was designed to bring about some state of affairs.

[15] See DEFINITION5.2 for a description of design by natural selection.

References

Åqvist, L. [1984], "Deontic Logic", *Handbook of Philosophical Logic*, Volume II: *Extensions of Classical Logic*. Ed. Gabby, D., and Guenthner, F., Boston, MA: D. Reidel Publishing Co.

Bull, R., & Segerberg, K., [1984], "Basic Modal Logic", *Handbook of Philosophical Logic*, Volume II: *Extensions of Classical Logic*. Ed. Gabby, D., and Guenthner, F., Boston, MA: D. Reidel Publishing Co.

Catach, L., [1988], "Normal Multimodal Logics", *Association for the Advancement of Artificial Intelligence Proceedings*. 88, pp. 491-495.

Fitting, M. and Mendelsohn, R., [1998], *First Order Modal Logic*. Boston, MA: Kluwer Academic Publishers.

Hughes, G., and Cresswell, M., [1996], *A New Introduction to Modal Logic*, London: Routledge.

Kripke, S., [1963], "Semantical Analysis of Modal Logic I, Normal Propositional Calculi", *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9, pp. 67-96.