

# The dialectic of life

Christopher Shields

Received: 24 June 2010 / Accepted: 6 January 2011  
© Springer Science+Business Media B.V. 2011

**Abstract** In the dialectic of debates about the extension of life, one witnesses a predictably repeating pattern: one side appeals to a motley of variegated criteria for something's qualifying as a living system, only to find an opposite side taking issue with the individual necessity or collective sufficiency of the proposed criteria. Some of these criteria tend to cluster with one another, while others do not: metabolism, growth and reproduction; self-organization and homeostasis; an ability to decrease internal entropy by the appropriation of free energy; stimulus response suited to self-preservation and propagation; and adaptation. In competing approaches to the extension of life, these sundry criteria thus jockey for authority, with one group of theorists promoting some subset of them as essential for life, where the appeal to essence is as likely as not to be a simple modal *sine qua non*, and another denying that the nominated criteria are really necessary at all. The debate then stalls, because there seems to be no shared methodology for adjudicating such disputes. We may address this unhappy situation successfully by coming to appreciate that life is what we may call a core-dependent homonym.

**Keywords** Life · Univocity · Core-dependent homonymy · Family resemblance

## 1 The extension of life

In the mid-twentieth century, two philosophers quarreled about a rosebush, or, if you prefer, about a fake rosebush.<sup>1</sup> Confronted with the case of a rose garden filled with all

---

<sup>1</sup> These opening paragraphs reproduce, with certain liberties, a debate between Putnam and Ziff, as characterized in Putnam (1964). See also Matthews (1977).

---

C. Shields (✉)  
University of Oxford, Oxford, UK  
e-mail: christopher.shields@lmh.ox.ac.uk

normal rosebushes save one, they differed on whether the peculiar instance qualified as a living being. The exception to the norm looked very much like all the others, grew larger and smaller as they did, and even gave off a lovely damask scent, but upon closer inspection was revealed to be not a rose grown from a seed but rather an ingenious device constructed artificially: It was controlled from within by a series of tiny levers and pulleys, and “bloomed” along side the others only because it had been coded to do so according to a pre-timed program. One philosopher was clear in his mind that the exceptional case was an inorganic artefact, and so not a living being at all, and so trivially not a rosebush either but rather an impressive facsimile thereof. The other philosopher was not so sure: At any rate, its being inorganic was not by itself an impediment to its being alive. It seemed to him parochial to insist that as a matter of metaphysical necessity only the organic might manage to live.

These philosophers were arguing inter alia about the extension of *life*, about the range of living beings and about the requisites of falling into that range. From a certain angle, it may appear that their dispute was largely idle, that there was nothing really to decide the issue dividing them, that they were really quibbling about the words they used and nothing more. One of them, it seems, would have been content with that resolution, and even wrote a book called *Semantic Analysis*<sup>2</sup> intended in part to put to rest philosophical disputes by focusing attention on entrenched patterns of linguistic habit, implying that there was no further profitable kind of analysis available to philosophers. The other was again not so sure: Could not life be multiply realisable, such that we might in principle develop a higher-order account at least consistent with the hypothesis that something inorganic might be alive? After all, as far as linguistic habit is concerned, and as a purely descriptive matter, we find competent adult members of natural language communities referring to all manner of subjects as alive: gods, angels, science-fiction robots, real-life aliens, if there are any, and even, in the guise of the Gaia Hypothesis, the Earth itself. If such beings are, or are possibly, within the extension of life, then we have no immediate cause to disbar the mechanical rosebush, at least not as a consequence of its being inorganic.

Matters have only become very much more complicated in the half-century since those two philosophers conducted their inconclusive dispute. Consider first the occult-sounding Gaia Hypothesis—initially, it is worth recalling in view of some of its more colourful exponents, introduced by a NASA-engaged scientist as an empirically motivated hypothesis.<sup>3</sup> The proponents of this hypothesis advance the world eco-system as sufficiently homeostatic, cybernetic, and self-regulating to qualify as a single living organism. Needless to say, the Gaia Hypothesis has met with all manner of detractors, including not least those who allege that it is implicitly teleological and so rests on an outmoded framework of explanation. Further, according to some of these same detractors, the world as a whole obviously has no DNA distinct from the DNA encoded in the carbon-based matter of the uncountably many animals, plants, bacteria, fungi, and assorted archaea that populate it. Such beings qualify as living systems not because

---

<sup>2</sup> Ziff (1960).

<sup>3</sup> For the original, scientific formulations of the Gaia Hypothesis, see Lovelock and Margulis (1974) and Lovelock (1990, 2000).

they achieve a degree of unity dependent upon an end-state or *telos*, but rather because they exemplify empirically ascertainable patterns of DNA implementation and replication. This the world ecosystem lacks.

Here again the debate is about the extension of life. One party brooks no extravagant speculation: Life extends to human beings and other animals, to plants, whether cultivated or uncultivated, and finally to the many micro-organisms recognized by cell- and micro-biologists. No further. The other party is prepared to countenance a broader range of contenders, thus contemplating a much broader extension for life. After all, they note, many people seem to suppose that angels exist as living beings, or at least that they *could* exist as a matter of logical possibility. In retort, the hardliners scoff, insisting that nothing lacking magnitude could be thought to metabolize, and given that metabolism is necessary to life, angels are worse than a fanciful fiction: They are positively incoherent.

An analogous dispute plays out back on Earth. Extremophile archaea and bacterial parasites qualify as living systems on most accounts, whereas acellular viruses divide the taxonomizers. Some treat them as living, on the grounds that they exhibit descent with modification, whereas others are content to exclude them in view of one of their more angelic traits: They do not metabolize. Moving from viruses into the contentious zone of artificial life, matters grow hazier still. The possibility of artificial life (ALife), whether weak or strong, and if strong whether soft, hard, or wet—that is, whether software-, hardware-, or biochemically-based—further flummoxes the field.<sup>4</sup> If we can create systems that do the things that indisputably living systems do, or which at least pass the most demanding Turing test for life we can fathom, then by what right would we bar such beings from the extension of life? Similarly, looking beyond these earthbound disputes but still shy of the supernatural, there arises the possibility, or presumed possibility, of alien life forms, the quest for whose existence is taken sufficiently seriously in some quarters that it garners extravagant funding, despite the fact that Astrobiology is as yet a discipline without a domain.

In the dialectic of these debates about the extension of life, one can witness a repeating pattern: One side appeals to a motley of variegated criteria, only to find the opposite side taking issue with their necessity. Some of these criteria tend to cluster with one another, while others do not: metabolism, growth and reproduction; self-organization and homeostasis; an ability to decrease internal entropy by the appropriation of free energy; stimulus response suited to self-preservation and propagation; and adaptation. In competing approaches to the extension of life, these sundry criteria jockey for authority, with one group of theorists recommending some one or subset of them as *essential* for life, where the appeal to essence is as likely as not to be a simple modal *sine qua non*, and another denying that the nominated criterion is really necessary at all. The debate then stalls, because there seems to be no shared methodology for adjudicating disagreements about which criteria, if any, really are necessary.

There is further, and more revealingly, a vertical axis along which many of these exchanges may be fruitfully compared. Some of our theorists, whom we may call

---

<sup>4</sup> Weak ALife seeks only to create systems which model life; strong ALife contends that it is possible to create systems that do not merely represent or model life, but actually live.

*organicists*, see life as biologically or organically enmeshed, contending that all life is DNA-encoded, with the result that only physical systems with appropriate chemical compositions qualify as living. Others recoil, accusing the organicists of suffering from a small sample problem. The living systems in our immediate neighbourhood may all look a certain way, but our neighbourhood is surely, by any account, puny. At the opposite, top end of the vertical spectrum are *abstractivists*, who are inclined to think of living systems as DNA- and even carbon-independent. Those with abstractionist tendencies find an able champion in the Hungarian mathematician and physicist John von Neumann, who maintains, simply, that “life is a process which can be abstracted away from any particular medium.” The emphasis on *process* is key to such feats of abstraction: While DNA implementation from an organicist perspective requires a polymer made of simple nucleotides, from an abstractivist perspective what implements DNA is effectively any suitable bearer of coded information capable of transcribing itself in receptive media. From the standpoint of an analysis of life, the abstractivist will thus happily set aside as a perfectly respectable but ultimately empirical matter the question of *which* implementation bases we find encoding DNA in our local area, and even, in more extreme cases, the necessity of life’s requiring DNA-encoding at all.

In the face of such methodological divergence, it has been tempting to settle the question of the extension of life with a resigned refrain: “It depends on what you mean by ‘life’.” It depends, that is, on what definition of the word ‘life’ we have antecedently adopted. If we are organicists, theses like the Gaia Hypothesis emerge as non-starters, along with the quaint suggestion that may be ethereal angels or the more mundane contention that we may come to countenance mechanically constructed living rose-bushes. Yet if we have opted for abstractivism, then we can at least entertain such hypotheses as possibilities, to be judged on the basis of whatever confirming evidence their devotees find themselves able to muster.

As a response to the untidy set of disputes about the extension of life, the it-depends-upon-what-you-mean-by-‘life’ refrain is both dispiriting and defeatist. It is dispiriting because it is too readily regarded as the final word on the matter, and it is defeatist because it presumes implicitly, or very nearly explicitly, that there is no fact of the matter about the nature of life. All that can be said about life can be stipulated; and beyond that stipulation there is no further worthwhile question left to answer. Perhaps this is because we are to assume that it is already known that life has no nature, or, more mildly, that even if there were a determinate nature to life, it would be beyond our ken, not least because it could not be an empirically determinable affair. In short, if the matter is not empirically decidable, and that is all the determination we can hope to achieve, then it will be otiose at best to attempt an account of life. Accordingly, from this perspective, there is no substantive issue about the nature of life to be joined. The reason that we have disputes about this question is simple: We have failed to come to grips with the fact that there is really no issue before us. It all depends upon what you mean by ‘life.’

The posture is further dispiriting because it is incorrect, or at best unacceptably pre-emptory. Clearly enough, our decisions about the extension of life are conditioned by our antecedent views about the character of life and living systems. Still, that by itself gives us no reason to suppose that decisions about the extension of life are ultimately

a matter of linguistic stipulation.<sup>5</sup> For instance, no one with a serious stake in this discussion should allow that it is utterly arbitrary which linguistic definition of life, if any, we are to adopt. No one will take seriously, for example, the suggestion that by ‘life’ is meant ‘all and only chrome-plated metal,’ so that the door handle of a 1963 Jaguar is alive but its antique-car enthusiast owner is not. For we do have at least some settled views about the extension of life.

Minimally, what might be acceptably contended at this juncture is that within some range of acceptable alternatives, presumably a range delimited by some or all of the criteria already mooted, there is no further fact of the matter as to which account of life is uniquely correct—there being, indeed, no uniquely correct account in the offing. Accordingly, one may suppose, the only question remaining concerns which among the various extension-setting criteria introduced ought to be favoured *for which purposes*. That is, if the question is not arbitrary, then neither is it more than pragmatically determinable.

Why, though, should we concede that even this conclusion is anything other than pre-emptory? The pragmatic impulse is rooted partly in a negative thesis, to wit *that there is no meaningful interest-independent analysis of life possible*. Where, however, is the argument for this conclusion? If it is only that we have failed thus far to achieve a *consensus omnium*, then it is plainly a *non sequitur*. If there is a principled argument available, then those preaching defeat need at least to state it. So far, at least, the arguments one encounters are redolent of what should by now be a moribund positivism.

Consequently, I contend, both the it-depends-upon-what-you-mean-by-‘life’ refrain and its pragmatic cousin are at this stage of our inquiry wholly premature. Moreover, to repeat, these negative conclusions, despite the self-mollifying tones of their issuers, are hardly free of charge: If it is to be shown, and not merely asserted, that no viable account of life is available to us, then this will require the production of some motivating reasons in the form of an argument. One thing that is wanted would be a precise account of what it is that cannot be done. Is the point that it is impossible to offer an essence-specifying analysis of life? Any empirically verifiable essence-specifying analysis? Any empirically verifiable analysis? Any account not at least implicitly relying on a theoretically introduced term?

These are differing sorts of negative findings, requiring different sorts of motivating arguments. In the absence of such arguments, it is salutary to focus on a revealing shortcoming in the current dialectic of life: There is a strong tendency, exhibited equally within the positive and negative sides of the debate, to accept a paradigm which is

<sup>5</sup> There is a tendency to suppose that life is either a natural kind or else merely a matter of linguistic preference. For a clear instance this pervasive attitude, see Cleland and Chyba (2002), who maintain that our lacking an account of life results precisely from our inability to develop a framework for treating it as a natural kind: “Perhaps life is not a natural kind. If it is not, how we define it will forever remain a matter of no more than linguistic choice” (p. 391). This is, however, a false dichotomy. If linguistic choice is constrained by data, whether extensional or intensional, then a definition of life is more than a *mere* linguistic choice, and what Cleland and Chyba contend is false. If, by contrast, linguistic choice is unconstrained by data, then, again unacceptably, we are free to call whatever we wish alive (say the west pediment of the Temple of Zeus at Olympia), and there is no further discussion to be had. Either way, there is no reason in advance to accept this conclusion, even absent the assumption that life is a natural kind.

unsuited to this domain of inquiry. Along our vertical axis, for example, both hard-line organicists and comparatively catholic abstractivists have supposed that a defensible account of life, if there is one, would need to be *univocal*—that is, that any defensible account would need to state the necessary and sufficient conditions for life in terms which would determine decisively for any putative candidate whether it falls within or without its extension.

I contend that this univocity assumption should be rejected and replaced with a more accommodating, but still perfectly principled framework. The sort of approach I have in mind has an ancient lineage, but has unfortunately fallen out of view. This approach derives from a definitional technique developed initially by Aristotle,<sup>6</sup> but can be suitably modified for a modern sensibility. Life, I argue, exhibits a striking order in multiplicity as what we may call a *core-dependent homonym*. Imprudent as it would be to contend that this approach settles all the issues we might wish to consider about the nature of life, it does at least offer an account that serves to stave off the easy and indulgent presumption that there is no fact of the matter for us to discover about nature of life. For this reason alone, I contend, core-dependent homonymy deserves a place at the table in the current dialectic of life.

## 2 Three approaches to the study of life: positioning core-dependent homonymy

Suppose we have not given up all hope for developing, or discovering, an account of life. If we remain optimists on this score, then we might at least initially be inclined to suppose that we have but one of two choices.

We might first, and most ambitiously, hold out for *univocity*. That is, we might hope for a non-disjunctive account of life answering to the following form:

- Life =<sub>df</sub> ...

Here the idea would be that we might yet uncover for life something akin to:

- Water =<sub>df</sub> H<sub>2</sub>O

This sort of definition is more than nominal. Instead, it is non-disjunctive and essence-specifying; according to its proponents, it captures the nature of water.

It is more than nominal, because it says more than what we might rightly say of water in a lexical definition:

- Water is the clear, potable liquid of which lakes and rivers are composed and which falls as rain and issues from springs.

This nominal definition, paraphrased from the *Oxford English Dictionary*, is correct as far as it goes; but it does not go far enough, since (the philosophers will tell us)<sup>7</sup>

<sup>6</sup> Shields (1999) develops and defends Aristotle's notion of core-dependent homonymy as it is advanced within the context of Aristotle's own theorizing. The current paper is an appropriation and development of that general approach to philosophical analysis, rather than a presentation or defense of Aristotle's own views.

<sup>7</sup> Putnam (1975/1985).

we can easily imagine a clear potable liquid which fills the rivers and lakes but which is not water—which is, say, a liquid composed not of hydrogen and oxygen but rather some hitherto undiscovered stuffs:  $x$ ,  $y$ , and  $z$ . If that is so, then the nominal definition hovers too close to the surface and so fails to capture the essence of the analysandum.

Further, the scientifically respectable definition of water is non-disjunctive. It does not tell us, for instance, that:

- Water =<sub>df</sub> (i)  $H_2O$  ; or (ii)  $x$ ,  $y$ , and  $z$ ; or (iii) ...

The worries about such disjunctive definitions are several,<sup>8</sup> but they include centrally the role of the ellipsis after (iii). If we do not know how many disjuncts we may yet uncover, then our account is as yet inherently and unacceptably open-ended. In that case, then we have not even specified the extension of *water*. (Compare: triangle =<sub>df</sub> a closed plane rectilinear figure having three sides and three angles). On the other hand, if we do know, in principle, how many disjuncts we will find, then we also know, evidently, that feature which serves to bring our definition to its close—and presumably *that* was the feature we were seeking all along.

Finally, a definition of this style, to bring into focus the way in which it is more than merely nominal, is *essence-specifying*. It is meant to capture the nature of water, by telling us not only how water *seems*, but how it *is*. It should, then, capture that feature without which nothing could be water, and with which anything most definitely is water. That feature would seem to be *being  $H_2O$* . Further, if we are hoping for more than a modal account of essence,<sup>9</sup> such a definition is welcome because it displays something deeply explanatory of water's behaviour and various nomological enmeshments.

As I have already suggested, most who recoil from the prospects of providing an account of life adequate to the task of settling at least the extensional question have some version of the univocal analysis in mind. It seems to many, that is, that there simply is no *one*, non-disjunctive essence-specifying definition of life available. So, they conclude, no positive account is forthcoming. That is, however, much too quick an inference, and for two distinct reasons, only one which has of late been recognized. The first, currently more common reason is derived from so-called *family resemblance* hypothesis inspired by Wittgenstein. As against what he regarded as the endemic philosophical disease of “craving for generality” where none is to be had, Wittgenstein counseled that we should come to appreciate that core philosophical concepts behave more like *game* than like, for instance, *triangle*. There is no one essence to all and only activities rightly called games. Still, Olympic games, board games, football games, and games of chance have over-lapping connections and resemblances, not unlike, in Wittgenstein's view, the kinds of physical connections and resemblances displayed by the various members of a large family. So, for instance, not all the Wilsons have exactly the same nose, high forehead and curly auburn hair, but they are all nonetheless recognizably Wilsons. Games, like the Wilsons, exhibit “a complicated network of similarities, overlapping and criss-crossing” sufficient to mark them all as belonging to the same general group, without its being the case that there is some one

<sup>8</sup> See Kim (1978).

<sup>9</sup> Fine (1994) successfully defends a non-modal account of essence.

general feature which they all sport.<sup>10</sup> So, although it is not arbitrary that all games are called games, nor is that due to there being some one essence shared by all and only games.

In the minds of many the suggestion lies near that what holds for games holds too for *beautiful* and *just* things, for *causes* and *laws*, and for *rights* and *wrongs*—and, to return to the case at hand, to *life*. The notion of family resemblance may indeed seem just right for life. After all, as we have seen, its extension is in dispute; new forms may yet be discovered or indeed developed; and no single property, even upon deep and honest reflection, seems to qualify as just that one feature which all living systems necessarily, as a matter of essence, manifest. If someone wishes to say both that angels and archaea are alive, then that is because these variegated systems somehow exhibit a sufficient number of the loose family of traits seen manifested by other, indisputably living systems. No more than that need be said; indeed, on this approach to definition, no more *can* be said. To say more is to indulge the deep but doomed desire for generality and univocity where the world supplies variegation and multiplicity. So, although it may prove illuminating to reflect upon these sundry traits of life and to catalogue their relative frequency, there is nothing further to be uncovered about its nature. Living systems are as games: they exhibit family resemblances and no more.

Whatever its positive commitments,<sup>11</sup> the doctrine of family resemblance has a lot in common with the negative stance presupposed in the it-depends-upon-what-you-mean-by-‘life’ refrain. According to its proponents, we can investigate the multifarious criss-crossing attributes of living systems secure in the knowledge that no univocal account is forthcoming. Unfortunately, the two stances also have something more in common on this score: neither offers a principled reason for advancing its negative verdict. Accordingly, before we can begin the task of detailing the attributes appearing in the fuzzy resemblance set of life, we require reassurance that no univocal account can yet be established.

More importantly, and this leads to our third and preferred style of analysis, the advocates of family resemblance may even be right about their (yet to be established) negative finding without having license to infer that the best we can hope for is family resemblance. This is because there is a *tertium quid* between univocity and family resemblance; and whether or not this position should be occupied, it is undeniable that univocity and family resemblance do not exhaust the logical space of positive definition. Between them is the once popular but now mainly neglected Aristotelian alternative of *core-dependent homonymy*. *Homonyms*, in the sense here discussed, are entities with common names but diverging accounts. Thus, banks—the sides of rivers and the institutions where money is kept and traded—are homonyms. These sorts of homonyms are, of course, philosophically uninteresting.

Far more interesting are *connected homonyms*, as for example, arguably the kinds of *conscious* beings marked, for instance, by Rosenthal’s distinction between *creature* and *state* consciousness.<sup>12</sup> Some sorts of things, including human beings, are

<sup>10</sup> *Philosophical Investigations* §§165–166.

<sup>11</sup> For some caution about the positive force of Wittgenstein’s notion of family resemblance, see Sluga (2006).

<sup>12</sup> Rosenthal (1986).



conscious; others, including the states and processes of human beings, are also called conscious, though now in a different but related sense. That there are two senses in play can be appreciated by a technique developed by Aristotle, in his *Topics* i 15: the test of opposites. A being that fails to be conscious in the creature sense, say a stone or a sofa, is not unconscious, but rather simply *non-conscious*. By contrast, something that fails to be conscious in the state sense, perhaps a desire or a seething resentment, is not *non-conscious* but *unconscious*. Philosophers have further distinguished kinds of state consciousness. Thus, for instance, Block cautions that we must keep separate *phenomenal* and *access* consciousness, roughly the distinction between the qualitative character of mental states and those cognitive or representational features which are poised to control and direct behaviour.

These distinctions are sometimes disputed and are in fact variously formulated; but that is not our current concern. Rather, we may in the present context note two things about such distinctions among kinds of consciousness: They are sufficiently close, according to their proponents, that they escape the notice of many; and, yet, on the assumptions that the distinctions are in fact genuine, they stand in the way of any univocal account of consciousness. In Block's colourful language, consciousness is a *mongrel concept*.<sup>13</sup> We may prefer to say, for reasons to be clear presently, that, again on the assumption that these distinctions are legitimate, conscious things are *connected homonyms*. The sorts of distinctions in play are subtle ones, so subtle indeed that even their proponents recognize that they require argumentation and special pleading. Even so, there remain important connections between the various forms of consciousness in play. Surely, at a bare minimum, these ranges of phenomena are not mere (i.e., bank-style) homonyms.

What are here called "connected homonyms" may so far, however, appear to be neither more nor less family-resemblance kinds called by a more arcane name. That complaint is indeed apt as far as it goes, but it fails to appreciate where further we might wish to go in our thinking about connected homonyms. Certain important cases of connected homonymy are not merely connected, but are rather connected in systematic ways, ways which enable us to see much more order than we find exhibited in family resemblance kinds. To appreciate this, we may begin with a simple illustration, one of Aristotle's favourites. Consider the following sentences:

1. Socrates is healthy.
2. Socrates' complexion is healthy.
3. Socrates' diet is healthy.
4. Socrates' decision to take up jogging is healthy.

For ease of explication, we may speak in a linguistic idiom, but only on the provision that we bear in mind that we are in the end not thinking about the lexical definition of the word 'life,' but rather about the question of whether we may yet uncover a deeper, non-nominal definition of *life*. With that proviso in mind, we may note two things about the predicate '...is healthy' in (1)–(4). First, they do not mean the same, and, second, their meanings are related, so that we do not have bank-style homonymy here.

<sup>13</sup> Block (1997, p. 376).

Regarding the first point, while complexions, diets, and decisions may all be called healthy, what is meant in each case differs. This is shown by an easy paraphrase:

1. Socrates is flourishing in body and mind.
2. Socrates' complexion is *indicative of his being healthy*.
3. Socrates' diet *helps to maintain his health*.
4. Socrates' decision to take up jogging *will, if enacted, contribute to his health*.

Still, and this is the second point, while they establish some difference in meaning, these paraphrases all reveal that the predicate '...is healthy' as it shows up in (1)–(4) are closely connected, and in a specific way. Of course both the connections and the differences in play in this example are more obvious than the putative connections between the various forms of consciousness already introduced, but that only because the predicate '...is healthy' serves as an especially useful illustration of core-dependent homonymy.

That is, there is a significant order obtaining in the predicates in play in (1)–(4); and it is here that the notion of core-dependence leaves behind family resemblance and connected homonymy. For there is more than mere connectedness: The predicate '...is healthy' as it occurs in (2)–(4) bears an asymmetric dependence relation to its occurrence in (1). Or, more exactly, accounts of 'healthy' as it occurs in (2)–(4) perforce make reference to an account of its occurrence in (1), though the opposite does not obtain. Thus, for instance, to explicate what it means to say that "Socrates' complexion is healthy" one must advert to *health*, but not conversely. Since Socrates' complexion is healthy because, and only because, complexions of his sort tend be indicative of health, any explanation of the suitability of the predicate in its derivative application must reflect that fact. There is, thus, account dependence: Non-core instances of the predicate require mention of the core instance in the specification of their accounts. Further, again, this account dependence is asymmetric. Indeed, even though there is likely a perfectly good nomic connection between health as it occurs in human beings and the healthiness of healthy human complexions, there is no need to explicate health in terms of the facts about complexions. So, in cases of core-dependent homonymy, we find a certain sort of structure in the multiplicity of cases: despite the non-univocity of the predicate, the internal relations between their various accounts reveal core-dependent order rather than mere family resemblance.

This order, slightly more formally, takes the following form:

*a* and *b* are homonymously F in a core-dependent way *iff*: (i) *a* is F; (ii) *b* is F; (iii) the accounts of F-ness in '*a* is F' and '*b* is F' do not completely overlap; and (iv) the account of F in '*b* is F' necessarily makes reference to the account of F in '*a* is F' in an asymmetrical way (or *vice versa*).

Given this apparatus, it follows, at a bare minimum, that there is logical space for an account of life that is neither univocal nor merely an instance of family resemblance. It further follows that if *this* is what we mean by 'life,' then we mean something which is neither arbitrary nor in any interesting way a matter of linguistic stipulation. Rather, by 'life' we mean something exhibiting an arresting order across its unending variety.

### 3 The core-dependent homonymy of life: a proposal

Is this what we mean by life? After all, one may fully agree with the point about core-dependent homonymy's logical space without also agreeing that life occupies it. In general, one may even allow that this space is in fact sometimes occupied, perhaps by *health*, without committing further to the particular applications to which its proponents have sometimes wished to put it. For example, the originator of core-dependent homonymy, Aristotle, immediately endeavored to extend the apparatus to such comparatively abstract notions as *cause*, *unity*, and even, most controversially, *existence*. He thereby met strong resistance, even from those favourably disposed to the style of analysis it commends.

So, what of life? We may make progress only by investigating at least one substantive proposal for life as a core-dependent-homonym.<sup>14</sup> As will be immediately plain, the account attempted inclines towards the abstractivist end of the spectrum:

- $x$  is alive  $\equiv_{\text{chdf}}$   $x$  is an intrinsic teleonomic system

The first thing to notice about this proposal is that non-standard definitional sign ' $\equiv_{\text{chdf}}$ ', where ' $\sigma \equiv_{\text{chdf}} \Phi$ ' is to be read as "the core of a core-homonymous definition of  $\sigma$  is  $\Phi$ ." This core is the feature to which every non-core instance must allude in its account, even though this core does not provide anything like necessary and sufficient conditions for all the entities called ' $\sigma$ .' Thus, to illustrate, if the core of health is, let us say, *flourishing*, then 'health' as it appears in 'her exercise regimen is healthy' needs to include the notion of flourishing in its account. So although there is no one univocal definition for all cases of health, every account of health needs ultimately to make reference to the notion of flourishing, since every account of a non-core case needs to cite the core instance of health in its specification. Thus, in the current illustration we have: 'Her exercise regimen is one which *promotes her health*—that is, her flourishing.' We would not wish to say, however, that her regimen is itself flourishing. With some complications yet to be noted, the same obtains for *life*.

In order to assess the plausibility of this proposal, we may proceed simultaneously on two levels. First, we may consider the actual core-homonymous definition I wish to propose. Second, we may simultaneously reflect more abstractly on the form of definition this account takes. It is worth proceeding on both levels at once, since one may dispute the terms of the definition offered even while coming to appreciate the relative advantages of the core-dependent homonyms over their main competitors, which I take to be univocity, family resemblance, and the it-depends-upon-what-you-mean-by-'life' refrain. This more abstract level is, so to speak, up one level from the material level of assessing the core-homonymous definition proposed. By reflecting on this higher order matter as we proceed we can assay more abstractly the question of whether core-dependent homonymy is the appropriate framework for thinking about life.

<sup>14</sup> Although I myself accept this proposal, and introduce it as a serious contender, I do not in the current paper mean to argue for it extensively. For the reasons given in § IV below, I am disinclined to think that any ultimately compelling univocal account of life given in teleological terms is forthcoming. My emphasis here is on the *style* of definition it captures, and I introduce it primarily to illustrate how in principle an analysis of life as a core-dependent homonym might be structured.

To the first task, then, of assessing the definition at the material level. Plainly two terms in the definition cry out for elucidation: *intrinsic* and *teleonomic*. This account treats the core of life, that feature to which any account of a living system must make reference, as involving systematicity of a teleonomic sort. Immediately, then, it will encounter the charge leveled against the Gaia Hypothesis, to the effect that it adverts to an outmoded form of causal explanation. Without engaging that complaint fully,<sup>15</sup> it serves to observe first that teleology is not in all quarters regarded as having been superseded by superior forms of explanation,<sup>16</sup> nor even thought to be in competition with them.<sup>17</sup> It seems, in any event, pre-emptory at this early stage to rule out the account on such grounds, without first seeing what sort of explanatory work it can do. That work, I suggest, is considerable.

Further, although the account is intended to be fully realist in its ascriptions of end-states that organize and co-ordinate disparate sub-systems around a single directionality, it need not be so construed. That is, various forms of pragmatists could accept this definition in the spirit of Dennett, to whose notion of *an intentional system* this account may be regarded as a realist counterpart. For Dennett, an intentional system is any system ‘whose behavior can be—at least sometimes—explained and predicted by relying on ascriptions to the system of beliefs and desires (and hopes, fears, intentions, hunches, ...).’<sup>18</sup> One could say the analogous thing about teleonomic systems, viz. that they are any systems whose behaviour can be explained and predicted by relying on ascriptions of end-directedness. So, without waxing Paleyesque, we might agree that if in our explorations of the dark side of the moon we stumbled upon a box with the configurational complexity of a Macintosh G4 computer we should assume as a working hypothesis that it was an artefact organized for one or more activities, and that it was not (even while conceding to the philosophers in the room that it *could be*) an entirely accidental being, produced at random by the aleatoric processes of the winds upon the surface of the moon. On the contrary, we would do well to adopt the teleonomic stance toward it. We would do well, that is, to think that its sundry components were organized in lawlike and explicable ways around some function. We would, for what it is worth, most likely be right about that.

Now, as suggested, although the account is not intended to embrace this kind of pragmatism, the possibility of doing so does help highlight the general flexibility of this style of definition. One might, for instance, deny that there are *really* such things as intrinsic teleonomic systems, while allowing that it sometimes behoves us to adopt an explanatory strategy imparting just such unity to certain collections of molecules. Be that as it may, the core-homonymous definition advanced above is in fact intended

<sup>15</sup> For two different accounts of function available to a teleological account of systematicity, see [Wright \(1973\)](#) and [Cummins \(1975\)](#). [Godfrey-Smith \(1993\)](#) argues with some success that these two competing accounts are not really distinct after all, and that they tend to converge around a single approach. In any event, the notion of function invoked in the current discussion is by design neutral between these alternatives. What matters in the current approach to life is only that *some* approach to non-designed function remains viable. See also [Neander \(1991\)](#) and [Wright \(1976\)](#).

<sup>16</sup> See [Hawthorne and Nolan \(2006\)](#) for a crisp account of how teleological causation might best be understood.

<sup>17</sup> [Enç \(1979\)](#).

<sup>18</sup> [Dennett \(1987\)](#).

to capture the widely held view that living systems, by and large, have privileged forms of unity. That is, to take a paradigmatic case, a living animal is a system moving about the Earth as a unified whole, made up various parts and subsystems interacting in finely tuned ways. One could, as various philosophers inspired by Hume do even today, regard all such “unity” as intention-dependent. One could hold, for instance, that a leopard is no more a unified entity than that “entity” composed of exactly the following parts: the flame of the Olympic torch at 2.21 p.m. Beijing time on 21 March 2008, the left half of *The Triumph of Death* by Pieter Brueghel the Elder, and a wistful glance of Alcibiades cast in the direction of Socrates in 411 B.C. There are those proponents of universal mereological aggregation who will insist that every aggregate is perfectly on par with every other, and that there are no privileged unities. One could join them in thinking this way and then fall back on an instrumentalist approach to teleonomic unity; but in general I assume that one should preserve privileged, non-intention-dependent unity if possible. In any event, however that issue is decided, we need not let the unappealing possibility that *all* unity is intention-dependent derail our account of life.

This brings us to *intrinsicity*. If we think, as most of us do, that living systems are unities of a certain sort—a privileged sort—then we need either to find some principle of unity or to accept their unity as basic and inexplicable. Whatever that principle may be, whether or not it is an end-state that various systems and sub-systems subservise, it appears that it must be intrinsic in at least the following sense. Among the kinds of unities we are prepared to recognize as being organized around a single function, some, but only some, belong to a category we have already mentioned, namely *artefacts*. An artefact, for instance, a computer, is plainly a unit bounded in space and time, discrete from entities in its environment and the other entities with which it is contiguous (the table on which the computer sits is not itself part of the computer). The source of an artefact’s unity normally consists in its being integrated around some single function (in this case, computing). Still, although the account welcomes rather than precludes the possibility of robotic or artificially created life, no one at present regards his laptop as alive. Rather, artefacts are teleonomic systems that are derived, rather than intrinsic.

Now, it should be admitted that it is difficult to say precisely what is meant by intrinsicity in this connection, but it is possible to make a start, at least negatively. First, we should not insist that a system fails to be an intrinsic teleonomic system simply because it was caused to come into being by another teleonomic system, whether intentionally or not. After all, if we otherwise believe that humans are intrinsic teleonomic systems, as I do, we should not be barred from doing so by the bare fact that humans have parents. This form of causal dependence is irrelevant to the question of whether something qualifies as a system with an intrinsic end. By contrast, if we think that the identity conditions of an object depend upon its being employed as the kind of end-directed system it was intended to be, then we have grounds for supposing that its end is not an intrinsic one. A computer, for instance, like other artefacts we currently employ, fails to be an intrinsic teleonomic system on these grounds. Still, nothing prevents its becoming the case that something created by a human intentional system should come to acquire an intrinsic end. Presumably, for instance, if a scientist synthesized a molecule-for-molecule duplicate of an uncontroversially living system,

say, a human being, then the duplicate would be alive no less than the original. In any event, the cost of denying that result would be high, since then we would need to allow, for no principled reason, that two otherwise qualitatively indistinguishable beings could differ only in respect to being alive. In such a case, however, we would seem to allow that something might intentionally be fashioned to be such that it has an intrinsic end. For what it is worth, this seems to me already to be the case with respect to various genetically engineered bacteria: I feel no impulse whatsoever to suppose that they are any less alive than any other bacteria; nor do they seem to me to be any less lacking in intrinsic directionality.

Perhaps this is because it is further the case—although this introduces an even more contentious matter—that intrinsicity carries with it a robust form of normativity. An intrinsic teleonomic system exhibits normativity in its own right and not merely parasitically on the good of another: Whatever has an intrinsic end functions well or poorly in a manner that is good or bad *for it*.<sup>19</sup> Among the forms of functioning systems we can imagine, success may in some cases be good *for the system whose function it is* and in others not. Arguably, though again the matter is disputed, if you destroy my computer in a pique of rage, you have not done something to harm *it* so much as me. In any event, it does not seem to have a grievance against you. I do. Admittedly, this distinction is hard to make precise, but if we are prepared to introduce even this much native normativity into the notion of intrinsicity, then an intrinsic teleonomic system will be one about which it makes ready sense to ask whether it is flourishing. This, though, brings us squarely into the province of life.<sup>20</sup>

Now, to stave off one misunderstanding: It causes no difficulty to the account that there will be grey cases, and cases where judgments of intrinsicity are difficult to make. On the contrary, the presence of grey between the white of the living and the black of the nonliving will presently be claimed as confirmatory of the approach to life here defended. Core-dependent homonymy positively predicts greyness; since the account of life introduced is an instance of core-dependent homonymy, we should expect and welcome greyness at the margins. Indeed, to complicate matters still further, it may well be that the notion of intrinsicity is itself homonymous, though so much is not required for the current account to meet the objection. It will suffice for now to note that difficulties in determining intrinsicity tend to confirm rather than undermine the suggestion that life is the sort of core-dependent homonym envisaged.

<sup>19</sup> For a clear expression and plausible defense of the normativity of the teleological, see Bedau (1992). As he notes, the suggestion that teleology is inherently value-laden is not new (he provides useful references in 781 n. 1). What is new is his nuanced defense of a value-based analysis, together with a compelling set of criticisms of attempts to squeeze the normativity out of teleology. Of course, there are those who will regard Bedau's success as a failure for teleology. The current paper adopts the opposite tack in a way which tends to corroborate Bedau's analysis: Since we can in fact distinguish the living from the non-living, and can do so only by appeal to an inherently normative teleology, then we have reason to accept the explanatory fecundity of this framework. See also Schrödinger (1945) and Van Inwagen (1990).

<sup>20</sup> One reason this distinction is hard to make precise is that artefacts have hitherto not in fact been intrinsic teleonomic systems. Nothing, however, precludes their becoming so. Much more needs to be said about the nature of artefacts in this debate, though at this juncture it bears stressing only that the account of life given is intended to be permit the possibility of hard ALife: nothing about something's being created by an intrinsic teleonomic system precludes its being an intrinsic teleonomic system in its own right. For more on the character of artefacts, see Hiipinen, R, 'Artefacts' (<http://plato.stanford.edu/entries/artefact/>).

With that noted, it becomes possible to consider some further features of our account of life as a core-dependent homonym.

The core-dependent account before us has at least three further advantages, all of which in various ways reflect both its material and the formal features: (i) plasticity; (ii) essence-specification; and, consequently, (iii) extensional adequacy, including modal and not merely factive application conditions.

First, the account is plastic, in the sense that it allows for multiple realisability and open-endedness. In the first instance, we should not wish to restrict life, on some sort of *a priori* grounds, to systems made of the sort of stuff that constitutes us. If beings from Trafalmore made of alien goop visit us and discourse with us about the meaning of life, then we should conclude that they meet the definition; we should hardly deny them the distinction of life on the parochial grounds that they are not made of flesh and blood. Minimally, if something alien emerges with the kinds of intelligence we see realised in living systems such as ourselves, then we would have good grounds—grounds as good as or even better than the mere passing of a Turing test for life—to count them among the living. The core-dependent account under view easily embraces this kind of plasticity. The same holds more generally of teleonomic systems that do not manifest intelligence. If we have grounds for thinking that a system metabolizes, that it appropriates ambient sources of nutrition for its own flourishing, then, again, and for the same reason, we will have grounds for regarding it as a living being.

Indeed, in view of its plasticity, the current core-dependent account locates the debate about non-material living systems where it should be. The issue is not really whether life tolerates Cartesianism, but whether the prospect of immaterial beings engaging in activities and interacting with more familiar systems is a metaphysically coherent one. Given the abstractivist character of the current account, the question only becomes whether it is coherent to suppose that such systems could *implement* teleonomic systematicity. So, for example, there seems to be a fair agreement about what angels are supposed to be, namely, intelligent immaterial beings—which is to say a certain kind of immaterial intrinsic teleonomic system. So far, there is no contradiction involved in the hypothesis. Next, however, comes the rather more demanding contention that it makes perfect sense to think about processes like thinking in the absence of implementation bases utterly alien to anything we have in fact ever encountered. The burden then falls where it should: on the proponent of the hypothesis. Looked at that way, the debate needs to proceed at a comparatively abstract level, that of determining what might be required to realize the end-directed processes in view. That is the appropriate pitch for the battle, if anyone wishes to engage the fight: The necessity in question is merely hypothetical, and not immediately logical.

Second, now focusing primarily on the formal features of the account, though flexible because suitably plastic, core-dependent homonymy does not sacrifice the contention that there is something essential to life. To be sure, some will count this as a vice rather than a virtue, given what they regard as the bad odour of essentialism. Be that as it may, essentialism does have the welcome advantage of putting some constraints on the extension of life. By contrast, we seem for the most part to arrive at family-resemblance accounts by abstracting from the agreed extension of a concept, or, in its more linguistic guises, from the agreed application class of some word or other; but if that is so, then in cases where we have considerable numbers of disputes,

the limitations of that approach, however attractive it may have seemed to various Wittgensteineans, comes quickly to the fore. We know a family resemblance obtains when we already know that members belong to the same family; but we know that only if we have already set the extension. In the domain of life we find ourselves with deeply contested questions about the extension, questions often attempted to be decided pre-emptorily by appeals to the (putatively) necessary or essential features of living systems. If we accept the core-dependent homonymy of life, then we can step back from these debates, which in any event tend to devolve into unproductive stalemates, in order to judge the merits of candidates for life by their closeness to its core.

Third, and finally, core-dependent homonymy permits us to make distinctions at the right sorts of joints. For example, those who insist that living systems must be like us in this or that dimension (carbon-based, organic, self-replicating, DNA-encoding, subject to evolutionary development...) are rightly criticized on methodological grounds: They rely upon what is—or certainly seems to be—a small sample for extrapolation. They do so insofar as they draw a line around a known class of living beings, focus on one or another reasonably deep feature of those beings, and then demand that the feature hold of all actual or even possible living beings. When they are then criticized for getting the extension of life wrong, or, with equal force to many, with getting the *possible* extension of life wrong, they have little recourse but to discount the putative counterexamples as aberrant, unimportant, or simply impossible. The method of core-dependent homonymy escapes these sorts of short-sighted distinctions, but does so without setting an inflexible set of necessary and sufficient defining conditions.

Taking all that together, then, to the extent that these advantages of the account resonate, then we may draw two cautious conclusions at this juncture. First, we may embrace the thought that the core of life involves the unifying presence of some single unifying directionality, against which degrees of flourishing or failing may be judged. Second, if the imputation of end-indexed norms to living systems strikes some as extravagant or otherwise objectionable, it remains the case that the style of definition brings with it the twin advantages of second-order flexibility and essence-like depth of analysis. If we wish to be serious about settling disputes about life, and if we have not already given up in dispirited dejection, then these remain features to be preferred.

#### 4 An unwanted success?

If we have come this far with the account of life as a core-dependent homonym, then we have probably also already appreciated that objections lurk on all sides. One that has sometimes been voiced is surprising in philosophy, that the account is too successful for its own good. An initially indulgent criticism, the complaint is that the account heads too far in the direction of co-ordination. Indeed, it travels so far in that direction that it captures something it pretends is impossible: univocity. Here the thought is that if the account works at all, then it simply devolves into a straightforwardly univocal definition, such that in the end all the apparatus of homonymy is idle. The account really holds just this:



- $x$  is alive  $\text{=}_{df}$   $x$  is an intrinsic teleonomic system

So stated, the account is neither homonymous nor even disjunctive. Really, according to this critic, the account should take the form of the more traditional apparatus of ' $\sigma \text{=}_{df} \Phi$ ' instead of introducing the more complex framework of ' $\sigma \text{=}_{chdf} \Phi$ '.

This objection may seem as flattering as it is unlikely, though some have been inclined to press it.<sup>21</sup> Basically, the thought is that what is doing the work in the account is not its appeal to homonymy, but rather the appeal to teleology. This, after all, is what purchases the account's plasticity. What is more, I have myself earlier complained that various other approaches to life were defeatist, and at best pre-emptory, in claiming that no univocal account of life could be forthcoming. To be fair, the account mounted shares with these views the conviction that life is non-univocal, and so it no less than they owes an argument for this negative conclusion.

As uncommon as it is in philosophy to face an allegation of unwanted success, it must be said that this helpful criticism fails by being overly optimistic. It claims victory for univocity too readily. Seeing why will help highlight one last feature of core-dependent homonymy: Non-univocity may be non-apparent. In the case of *health*, of course, the non-univocity is apparent; but that only serves to explain why it serves a useful sort of illustration. The case with *life* is closer to the case of *consciousness*, where philosophers have understandably felt the need to argue for non-univocity. The thought about life (as I would also say about consciousness, though that is another matter) is that we first need to *establish* non-univocity, but to do so while avoiding the Block-style inference that it is a mongrel concept.

Some progress can be made in this area if we first reflect on the fact that the extension of life comprises an arresting variety of systems, from the entirely simple to the extraordinarily complex. Moreover, living systems may be complex or simple along a variety of dimensions. Some, like humans, are capable of internal representation, whereas others, like a rosebush, are not. Other cases are permanently less clear. Some wonder, for example, whether paramecia have internal representations, whereas for others this is a non-starter.<sup>22</sup> From the perspective of core-dependent homonymy, these differing intuitions, and the disagreements to which they give rise, merely reflect the fact that different sorts of systems, with different sorts of ends, may all yet qualify as intrinsic teleonomic systems. When we canvass even the range of undisputed living beings, we find that we are relying on all manner of different sorts of activities when coming to form judgments of systematicity. This alone already provides an impediment to any form of contentful univocity. If we say that gambling strategies, management techniques, and procedures for whipping egg whites are all *systems*, then perhaps we have said something true, but we have yet to capture any explanatorily rich essence common to these various activities. In a similar way, if we grant that all living things are intrinsic teleonomic systems only to discover that we have yet to explain any particular life activity, then we have yet to provide a contentful univocal

<sup>21</sup> I encountered reactions of this sort from members of the audiences at The University of Virginia and Stanford University when earlier versions of the ideas contained in this paper were presented.

<sup>22</sup> Fodor (1986).

account. If that is so, then short of core-dependent homonymy, we are left with family resemblance or defeat.

The situation may be illustrated by analogy to another disputed instance of core-dependent homonymy, namely *goodness*. Consider the following sentences:

- Despite its flaws, capitalism is good.
- Mother Theresa is good.
- The chorus at the Royal Opera is consistently good, better in fact than the chorus of any other major opera company.
- For families with pets, leather sofas are good.

One may observe that the speaker of each sentence finds some cause for approbation. From there, one might be tempted to move directly to the conclusion that there is some one general feature which qualifies the univocal account of goodness, namely *being esteemed worthy of approbation*. To establish non-univocity in the face of this much commonality, it is possible to rely on a two-part paraphrase test. In the first phase, we supplant the predicate ‘good’ with an appropriate paraphrase, as follows:

- Despite its flaws, capitalism is a just and efficient socio-economic system.
- Mother Theresa is a paradigmatically moral person.
- The chorus at the Royal Opera always performs to a high aesthetic standard, higher in fact than the chorus of any other major opera company.
- Leather sofas are durable and easy to clean and suitable for families with pets.

In the second phase, we observe that these paraphrases are not interchangeable. Thus, we cannot say sensibly that Mother Theresa is a just and efficient socio-economic system or that sofas are paradigmatically moral persons. If this is so, one may conclude that ‘good’ is non-univocal across these cases, even though the non-univocity is not immediately apparent.

The same two-phase process obtains in the case of *life*. Consider some of the instances of life within the extension recognized by the core-homonymous definition offered:

- Pavlov the dog is alive.
- The archangel Gabriel is alive.
- A laccaria fungus is alive.
- A laboratory-created android is alive.
- An extremophile single-cell bacterium is alive.

It is true that each of these is held to be a teleonomic system of some sort, just as all the instances of good were esteemed as in some sense worthy of approbation. Yet it is equally true that though they are all intrinsic teleonomic systems, a specification of the life activities in these predicates shows important differences in essences of the kinds of systems in question. In the first phase, then:

- Pavlov the dog engages end-directed perceptual and nutritive activity—where the end in question involves self-preservation and propagation.
- The archangel Gabriel engages in end-directed contemplation—where the end in question involves the glorification of God.

- A laccaria fungus engages in end-directed symbiotic transference with tree roots—where the end in question involves delivering minerals and amino acids to the tree whilst receiving synthesized sugars in return.
- A laboratory-created android engages in end-directed perception and thought—where the end in question, let us stipulate, does not involve self-propagation.
- An extremophile single-cell bacterium engages in end-directed rudimentary energy transferences in zones inhospitable to most life—where the end in question involves self-preservation and propagation by asexual division.

If we assume, even by stipulations, that neither the angel nor the android is capable of reproduction, then we find ourselves treating living systems as engaging in widely different kinds of activities with significantly different ends in view. To hearken back to our earlier set of criteria, we find ourselves relying centrally on reproduction in qualifying some beings as living, even while we deny that trait to other living systems. The same may be said about various other criteria across this range.

The second phase is then obvious. We can sensibly neither say that an extremophile single-cell bacterium engages in contemplation or that an angel engages in symbiotic mineral transferences. If so, then the sense of ‘alive’ in the predicates across this diverse extension is revealed to be non-univocal. In the variegated living systems we encounter, in both reality and imagination, ‘is alive’ will admit of importantly different sorts of analyses. We may conclude, then, that any appeal to a commonality between them will be so abstract as to miss what is essential to their individual forms of life. In this sense, ‘is alive’ is like ‘is good.’ Neither, then, admits of an informative univocal analysis.

Suppose that is so. We are then left wondering whether we are not again pushed too far in the opposite direction. That much diversity is too much diversity to be captured around a single core. This contention is, however, disputable. For in each instance, an account of the paraphrased predicate will need to make reference to the notion of an intrinsic teleonomic system. Thus, for instance, if we are to differentiate the kinds of chemical transferences occurring in the case of extremophile bacteria from the myriad other kinds of chemical transferences involved in nonliving systems in the environments of those bacteria, then we will have again to appeal to some kind of principle of unity, where, it has been argued, that available principle involves ascription of an end-state to that process. The same ineliminable appeal crops up in every other case as well, and for the same reason: Living systems are unities. Notably, this holds true even in the symbiotic cases: A laccaria fungus counts as a living system, rather than as a nonliving part of a living system (as your teeth are part of you), because it manifests its own intrinsic teleology. That an individual fungus also reliably serves the states of a living system beyond itself does nothing to diminish this point.

Taking that all together, then, the variety we find in the extension of living systems is incompatible with the existence of any contentful, essence-specifying univocal definition for their entire range. Even so, living systems, as living systems, do exhibit predictable and patterned kinds of unity. This much unity is given in terms of the intrinsic teleology of all such systems. This, then, is the core of life, and also, then, the core of a core-homonymous account of life.

## 5 Conclusions

We cannot hope in these brief remarks to have addressed all concerns about the prospects for an account of the nature of life, nor even to have allayed completely those few criticisms expressly entertained. Still, we may lay claim to four provisional conclusions, if only with decreasing degrees of security.

First, we should be unimpressed by those who pronounce that no account of life can be given. Most obviously, this negative finding requires a motivating argument, and, as we have seen, its proponents tend merely to declaim rather than to argue. When reasons do come to the fore, they are often as not inductive, appealing to our (putative) collective failure thus far. These champions of defeat are to this extent unscientific in their demeanor. Less obviously, but more importantly, this negative finding presumes that any account of life would need to be univocal, given in terms of non-disjunctive necessary and sufficient conditions yielding crisp extensions and epistemic transparency.

This presumption—and this is our second result—is unsustainable. Before we fall into despair on this score, we need to consider the definitional possibility of ordering life as a core-dependent homonym. Minimally, it is clear upon even a little reflection on the structure of definition that we can see an array of possibilities, ranging from univocity at one end of the spectrum to mere family resemblance at the other. What needs to be emphasized, and in a certain sense re-learned, is that this *is* a spectrum, that there are forms of definitional unity between univocity and family resemblance, and that prominent among them is core-dependent homonymy.

Further, while the apparatus of core-dependent homonymy may, let us grant, be available as an abstract alternative, its application to life remains, to be sure, a pending matter. This framework does, however, seem ideally suited to life. In this area we are confronted with wrangles about the extension of life generated in large measure by unspoken and often undefended conceptions of the necessary or essential criteria for life. These criteria are diverse to be sure, but are not so motley as to resist unification. They include, again: metabolism, growth and reproduction; self-organization and homeostasis; an ability to decrease internal entropy by the appropriation of free energy; stimulus response suited to self-preservation and propagation; and adaptation. To the extent that these criteria may tend to converge without merging into a single non-disjunctive essence-specifying account, they lend themselves a treatment given in terms of neither univocity nor family resemblance. So much recommends the framework of core-dependent homonymy for life. This recommendation is, then, our third result.

The fourth and final result is the most contentful, and, so also, inevitably, the most controversial: These various criteria converge, without yielding univocity, around the core notion of life as an intrinsic teleonomic system. Whatever its demerits, this approach carries with it a welcome sensitivity to the plasticity of life; to its seemingly scalar character; to its open-endedness of extension; to its in-principle ability to be synthesized in the laboratory, whether of chemistry or computing; and to its incipient normativity. Although this last feature is likely to rankle those already disposed to sneer at the imputation of norms to nature, it remains true that it makes sense to ask the following question of every living being: Is it flourishing? It is difficult

grasp how this question should be permanently present in the absence of the kinds of norms, whatever their origin, against which appraisals of life may be tendered and debated.

**Acknowledgments** This paper has benefited in two ways from the work of Mark Bedau. First, two of his papers on life have influenced my thinking about: (i) the normativity of the kind of teleology appropriate to living systems (Bedau 1992); and (ii) one sense in which life might be a sort of core-dependent homonym (Bedau 1998)—though, to be sure, that is not how he frames the matter. More proximately, he subjected an earlier draft of this paper to a series of insightful and challenging criticisms. Although I am doubtful that I have met them all satisfactorily, I am confident that the paper has been improved by his criticisms. I am grateful as well to another referee, Meg Scharle, who offered a different sort of extremely useful response. Again, though I cannot say with confidence that I have met all her worries, I feel sure that the paper is clearer for trying. I thank also Cissie Fu and Graham Oddie for instructive discussion, as well as the audiences at several universities where the ideas of this paper were first aired: Stanford University, New York University, The University of Virginia, and Dartmouth College.

## References

- Bedau, M. (1992). Where's the good in teleology? *Philosophy and Phenomenological Research*, 52, 781–896.
- Bedau, M. (1998). Four puzzles about life. *Artificial Life*, 4, 125–140.
- Block, N. (1997). On a confusion about a function of consciousness. In N. Block, O. Flanagan, & G. Guzeldere (Eds.), *The nature of consciousness*. Cambridge: MIT Press.
- Cleland, C., & Chyba, C. (2002). Defining 'life'. *Origins of Life and Evolution of the Biosphere*, 32, 387–393.
- Cummins, R. (1975). Functional analysis. *Journal of Philosophy*, 72, 741–765.
- Dennett, D. (1987). True believers. In D. Dennett (Ed.), *The intentional stance*. Cambridge: MIT Press.
- Enç, B. (1979). Function attributions and functional explanations. *Philosophy of Science*, 1979, 343–365.
- Fine, K. (1994). Essence and modality. *Nous Supplement*, 8, 1–16.
- Fodor, J. (1986). Why paramecia don't have mental representations. In P. French, T. Uehling, & J. H. Wettstein (Eds.), *Midwest studies in philosophy*, Vol X. Minnesota: University of Minnesota Press.
- Godfrey-Smith, P. (1993). Functions: Consensus without unity. *Pacific Philosophical Quarterly*, 74, 196–208.
- Hawthorne, J. P., & Nolan, D. (2006). What would teleological causation be? In J. P. Hawthorne (Ed.), *Metaphysical essays* (pp. 265–283). Oxford: Oxford University Press.
- Hiplinen, R. (2004). Artefacts. Retrieved July 9, 2009, from <http://plato.stanford.edu/entries/artefact/>.
- Kim, J. (1978). Supervenience and nomological incommensurables. *American Philosophical Quarterly*, 15, 149–156.
- Lovelock, J. E. (1990). Hands up for the Gaia Hypothesis. *Nature*, 8, 100–102.
- Lovelock, J. E. (2000). *Gaia: A new look at life on Earth*. Oxford: Oxford University Press.
- Lovelock, J. E., & Margulis, L. (1974). Atmospheric homeostasis by and for the biosphere: The Gaia hypothesis. *Tellus*, 26, 2–10.
- Matthews, G. (1977). Consciousness and life. *Philosophy*, 52, 13–26.
- Neander, K. M. (1991). The teleological notion of "function". *Australasian Journal of Philosophy*, 69, 454–468.
- Putnam, H. (1964). Robots: Machines or artificially created life? *Journal of Philosophy*, 61, 668–691.
- Putnam, H. (1975). The meaning of 'meaning'. In K. Gunderson (Ed.), *Language, mind and knowledge: Minnesota studies in the philosophy of science* (pp. 131–193). University of Minnesota Press. Reprinted in H. Putnam, *Mind, language and reality* (pp. 215–271). Cambridge: Cambridge University Press).
- Rosenthal, D. (1986). Two concepts of consciousness. *Philosophical Studies*, 49, 329–359.
- Schrödinger, E. (1945). *What is life?* Cambridge: Cambridge University Press.
- Shields, C. (1999). *Order in multiplicity: Homonymy in the philosophy of Aristotle*. Oxford: Oxford University Press.

- 
- Sluga, H. (2006). Family resemblance. In M. Kober (Ed.), *Deepening our understanding of Wittgenstein: Grazer Philosophische Studien* (pp. 1–21). Netherlands: Editions Rodopi B.V.
- Van Inwagen, P. (1990). *Material beings*. Ithaca: Cornell University Press.
- Wright, L. (1973). Functions. *Philosophical Review*, 82, 139–168.
- Wright, L. (1976). *Teleological explanations*. Berkeley: University of California Press.
- Ziff, P. (1960). *Semantic analysis*. Ithaca: Cornell University Press.