International Journal of Astrobiology 15 (4): 277–289 (2016)
doi:10.1017/S1473550416000021

Life is hard: countering definitional pessimism concerning the definition of life

Kelly C. Smith

Departments of Philosophy & Religion and Biological Sciences, Clemson University, Hardin Hall, Clemson, SC 29634, USA
e-mail: kcs@clemson.edu

Abstract: Cleland and Chyba published a classic piece in 2002 that began a movement I call *definitional pessimism,* where it is argued that there is no point in attempting anything like a general definition of life. This paper offers a critical response to the pessimist position in general and the influential arguments offered by Cleland and her collaborators in particular. One such argument is that all definitions of life fall short of an ideal in which necessary and sufficient conditions produce unambiguous categorizations that dispose of all counterexamples. But this concept of definition is controversial within philosophy; a fact that greatly diminishes the force of the admonition that biologists should conform to such an ideal. Moreover, biology may well be fundamentally different from logic and the physical sciences from which this ideal is drawn, to the point where definitional conformity *misrepresents* biological reality. Another idea often pushed is that the prospects for definitional success concerning life are on a par with medieval alchemy's attempts to define matter – that is, doomed to fail for lack of a unifying scientific theory. But this comparison to alchemy is both historically inaccurate and unfair. Planetary science before the discovery of the first exoplanets offers a much better analogy, with much more optimistic conclusions. The pessimists also make much of the desirability of using microbes as models for any universal concept of life, from which they conclude that certain types of 'Darwinian' evolutionary definitions are inadequate. But this argument posits an unrealistic ideal, as no account of life can both be universal and do justice to the sorts of precise causal mechanisms microbes exemplify. The character of biology and the demand for universality in definitions of life thus probably accords better with functional rather than structural categories. The bottom line is that there is simply no viable *alternative,* either pragmatically or theoretically, to the pursuit of definitions. If nothing else, the empirical data the pessimists demand will be a very long time coming and scientists will of necessity continue to employ definitions of life in the interim. Chastising them for this will only drive their ideas underground where they can escape critical analysis, making the problems caused by problematic conceptions of life worse.

Received 19 October 2015, accepted 26 January 2016

Key words: alchemy, alchymy, Cleland, Darwinism, definition, life, natural kind, pessimism, universal biology.

An introduction to definitional pessimism

"*Too often, philosophers' contributions to these questions seem designed only to reduce the number of thoughts that people can have, by suggesting that they have no right to some conceptions that they have or think they have. But equally philosophy should be able to liberate, by suggesting to people that they really have a right to some conception, which has been condemned by a simple or restrictive notion of how we may reasonably think.*"
- Bernard Williams

The publication of a much-cited article by Cleland & Chyba (2002) marked the beginning of a cottage industry arguing against the advisability of defining life – a movement I call *definitional pessimism*. Curiously, the authors admit not only the importance of an adequate definition of life, but also the *inevitability* of scientists using this in certain situations:

"*As science makes progress towards understanding the origin of life on Earth, as laboratory experiments approach the synthesis of life (as measured by the criteria of some definitions), and as greater attention is focused on astrobiology and the search for life on Mars and Jupiter's moon Europa, the utility of a general definition grows. In particular, definitions of 'life' are explicit or implicit in any remote in situ search for extraterrestrial life.*" (Cleland & Chyba 2002, p. 387)

Nevertheless, they express grave reservations about the possibility of formulating anything like an acceptable general definition, given our current state of knowledge.

Cleland and her collaborators are certainly not the first or only ones to express such sentiments (see Pirie 1937; Keosian 1974; Chyba & Mcdonald 1995; Frey 2000; Machery 2012)[1], but they are by far the most prolific defenders of the position, with more than a dozen papers and two books in recent years expanding on the pessimistic theme (Cleland & Chyba 2002; Cleland 2004; Cleland & Copley 2005; Cleland 2006; Cleland & Chyba 2007; Cleland 2007; Davies, *et al.* 2009; Bedau &

[1] See also Mix 2015 for an assessment of this trend from a more optimistic point of view.

Cleland 2010; Cleland & Chyba 2010; Cleland 2012; Cleland 2013a, b, c; Cleland & Zerella 2013; Cleland forthcoming). As a philosopher, Cleland is also typically more detailed and explicit in her arguments than other commentators, which makes focusing on her work a useful technique to gain insight into the pessimist mindset.

The 2002 paper is a beautifully concise statement of the core pessimist position. It lays out two sorts of arguments in brief, one philosophical and one pragmatic. Philosophically speaking, Cleland and Chyba claim that what we want in an ideal definition is an unambiguous way to differentiate categories of objects, in this case living and non-living entities. They claim that only when we can specify necessary and sufficient conditions that draw unambiguous lines between the groups in question have we met the minimum criterion for a scientific definition. This requirement makes it possible to have confidence that we have delineated what philosophers call a *natural kind* – that is, a reality independent of human convention. Conversely, to the extent that our putative definitions collectively fall short of this ideal, they conclude that science is not yet ready to, in the words of Socrates, 'carve nature at the joints,' and thus we should abandon attempts to formulate a definition.

To illustrate this point, we are invited to consider the example of alchemy. Before we understood that water is really H_2O, there was no way to cleanly disambiguate it from other compounds like nitric acid. The idea is that we are in much the same position at present with respect to understanding life: we can string together superficial properties, but have no real understanding of the fundamental nature of the thing we are trying to define. Thus, the 'definitions' of life one finds in the literature are better at elucidating the biases of individual researchers than uncovering reality. As evidence for this sad state of affairs, Cleland and Chyba point out that all candidate definitions are subject to 'robust counterexamples,' which they claim would not be the case if the definitions correctly described a true natural kind. Definitions of life that focus on evolution are singled out for particular scrutiny in this regard, with the resulting conclusion that they cannot, in principle, accommodate two sorts of counterexamples:

1. Weird life forms like those we see in Dyson's (1985) double origin theory, which would be have to be classified as non-living because they are not 'Darwinian.'
2. Sterile organisms like mules, since they cannot reproduce and thus could not count as evolving entities.

As a final consideration, Cleland and Chyba point out that any evolutionary definition will be very difficult to operationalize, echoing Fleischaker's (1990) worry about how long we would have to observe a candidate system to verify whether it is in fact alive in an evolutionary sense.

I will address these arguments and their variations in some detail in what follows, but I want to foreshadow what is to come with an initial set of responses I suspect have occurred to many of those who read the 2002 paper:

1. It is not clear whether entities like the ones Dyson hypothesized ever existed and, even if they did, it does not follow that our failing to describe them as alive is a problem. After all, any account of the origin of life will have to draw the line between living and non-living entities somewhere.
2. To intimate that we should consider systems like Dyson's alive clearly implies some concept of what life *is*, which strikes a discordant note in an argument against the advisability of any attempt to define life.
3. Using intuitive examples of life (mules, etc.) to make the point that we cannot define life is highly problematic. If we are interested in determining what counts as alive in a *scientific* sense, then the mere fact that a proposed definition does not accord with intuitive, pre-scientific concepts should not carry much weight.
4. If the ultimate goal is to discover the ultimate reality hidden beneath observable properties, why is the ease with which we can *confirm* our account critically important? To be sure, it would be a major problem if elements of a definition were unobservable *in principle*, but here the worry is simply that operationalizing an evolutionary definition might be more time consuming than we would prefer.

In many ways, the subsequent articles by Cleland and her collaborators merely elaborate on the themes of this initial paper. In what follows, therefore, I will examine each of these arguments in more detail.

What is a definition?

"*Persistent demands for rigor, clarity and verifiability sometimes scare ideas out of existence before they have a chance to come to fruition.*"

- Bertalanffy and Rappaport

I must spend some time addressing the philosophical argument concerning the nature of definition, though I will try to sail on these deep waters without imposing too greatly on the philosophical abilities of my readers. This kind of argument is harder to counter than others as the concepts are so complex they have been the subject of philosophical debate for thousands of years. To the extent they are taken seriously by the scientific community, I suspect it is largely because they are assumed to be philosophically uncontroversial – it is a bold biologist indeed who is willing to contend with the entire philosophical community on their home turf. Yet the actual state of affairs is more complex.

I certainly allow that the concept of definition Cleland advocates is a traditional one in philosophy of language and logic, with a supporting literature far too vast to summarize here. Fortunately, however, I need only to establish a modest claim using modest evidence: this concept of definition is philosophically controversial. The debate is perhaps best captured in a recent exchange between Boyd (1991) and Hacking (1991). Hacking defends a view similar to the pessimists', arguing that a natural kind (at least on a certain idealization) should:

...be defined by a set of necessary and sufficient properties (relations, etc.) such that, (i) the possession of these properties is, as a matter of fact rather than of logic, indicative of a very large number of other methodologically interesting properties and

such that, (ii) these defining properties are natural rather than social properties. (Boyd 1991, p. 127)

So Hacking, like Cleland, believes a natural kind is an entity picked out by clear necessary and sufficient conditions that are not themselves influenced by folk or socially based conceptions. If you accept this view, then you will expect any adequate definition to meet with near universal acclaim and encounter few anomalies to complicate its boundaries, which explains why the existence of 'robust counterexamples' is so damning in the pessimists' eyes.

Boyd, on the other hand, gives eloquent voice to a growing number of philosophers who reject this traditional account, ultimately concluding that:

"*…whether one accepts a realist or an empiricist diagnosis of the theory dependence of method, there are very good reasons for extending the conception of natural kind to include all property-cluster or social kinds reference which play any significant role in induction or social explanation*." (Boyd 1991, p. 134)

Thus, for Boyd, scientific classification is a profoundly theory-dependent activity and natural kind definitions should not be seen as static, but revisable in the light of new knowledge. If you accept this account, a good definition can still encounter significant opposition in the form of counterexamples because what it does, at least in part, is reflect socially constructed concepts. At first blush, it might seem that such a view is giving up on the task of identifying the ultimate reality behind our folk understanding, but it may in fact be the best way to *capture* that reality. To illustrate this point, Boyd discusses the concept of a species:

"*The paradigm cases of natural kinds – biological species – are homeostatic cluster kinds. The appropriateness of any particular biological species for induction and explanation in biology depends upon the imperfectly shared and homeostatically related morphological, physiological and behavioral features which characterize its members…The necessary indeterminacy in extension of species terms is a consequence of evolutionary theory, as Darwin observed: speciation depends on the existence of populations which are intermediate between the parent species and the emerging one…Any "refinement" of classification which artificially eliminated the resulting indeterminacy in classification would obscure the central fact about heritable variations in phenotype upon which biological evolution depends and would be scientifically inappropriate and misleading*." (Boyd 1991, p. 142)

In other words, any attempt to create an orderly categorization of inherently messy biological categories by imposing a rigid concept of natural kinds will *misrepresent* reality. Evolution may regularly produce the kinds of categories that fail to satisfy our psychological needs but, as the saying goes, 'such is life.'

Cleland responds directly to Boyd's approach in another article, saying that he:

"*…is thus using the term 'definition' in a nonstandard way. The important point for our purposes is that the pertinent homeostatic property clusters are not determined solely by means of an analysis of the concept that we associate with a natural kind term, and thus the use of the term 'definition' is misleading*." (Cleland 2012, p. 136)

Taking another philosopher to task for using terms in a non-traditional way is justified to some extent, since philosophers should at least pay homage to their shared conceptual history. However, Boyd can easily respond by arguing that the point he is trying to make is ultimately that the standard account of definition is problematic and thus we should *change* the way we talk about definition. To respond to this by complaining that the argument does not use the standard notion of definition runs the real risk of begging the question[2]. In any event, however one views the state of play within philosophy on these points, the fact that there is significant controversy among philosophers greatly weakens the force of the prescriptive claim that biology *must* conform to a particular philosophical account of definition.

Biology is different from physics

"*Nothing in biology makes sense except in the light of evolution*."

- Dobzhansky

Tension between the models of science put forward by philosophers and the practices of actual biologists is nothing new. Traditionally, philosophy of science has focused primarily on physics as the ideal of scientific practice (Hempel 1966, Cartwright 1980). The requirement that any adequate definition must include clean necessary and sufficient conditions for membership in a category may not seem unduly strict when discussing electrons (though even here there are complications), but it is been argued that biology is fundamentally different in this regard (Mayr 2007; Rosenberg & McShea 2007). That is, following Boyd's suggestion, biological categories like species may be messier *in principle* than the stock examples from logic and the physical sciences lead us to expect.

Let us revisit the concept of a species. Species are absolutely central to biological science, yet ever since Darwin, people have debated just what a species actually *is* (Mishler & Donoghue 1982; Ghiselin 1987; Mayden 1997). The situation is similar to the one we see with definitions of life – every putative definition has counterexamples and there are various camps, each pushing its own view. For example, the traditional *biological species concept* stipulates that a species must be reproductively isolated (Donoghue 1985; Mayr 2000; Noor 2002). This is still probably the most popular notion among biologists in general, despite the fact that *most* recognized species do not meet this requirement. Yet biology functions just fine despite this definitional heterogeneity – indeed, debates about the proper definition of a species have driven research that has greatly enriched biological theory, as with

[2] There are times when it seems the pessimists are simply debating who has the right to use the word 'definition' to describe their activities, but to the extent this is merely a semantic debate, it is not worth much attention.

https://doi.org/10.1017/S1473550416000021 Published online by Cambridge University Press

investigations into the possibility of (sympatric) speciation in the face of gene flow (Dieckmann & Doebeli 1999; Via 2001).

Of course, one could always conclude that, since biology fails to live up to our philosophical ideals, it is not really a science. Certainly much ink has been spilled in philosophy of science bemoaning the lack of truly universal laws in biology, a fact often taken to cast doubt on its scientific bona fides (Cooper 1996; Elgin 2006). Fortunately, this kind of conceptual imperialism has fallen out of favour in recent years as we recovered from the hangover of logical positivism and philosophy of biology came into its own. Most philosophers of science now seek to modify their philosophical conception of science to fit biology rather than the other way around.

So are there good reasons to think that biological explanation is different from explanation in the physical sciences? In a word, yes. There are two basic reasons for this, the first having to do with the complexity of the phenomena being explained. Relative to biology, the definitional task facing physics is simple: there are a manageable number of fundamental physical particles that interact according to a manageable number of fundamental physical laws to produce a manageable suite of behaviours in, say, an atom. However, as we progress up the levels of organization, through chemistry to biology and beyond, scientists are forced to deal with larger and more diverse casts of characters, possessing more and more degrees of freedom. As the number of interactive permutations grows, the boundaries between the categories they create also begin to blur[3]. So while there is (at least arguably) only one way to 'be an electron' there are many, many more ways to 'be alive,' or to 'be a dog.'

But there are also reasons to think the fuzzy boundaries between biological categories are not *simply* the result of complexity. As Dobzhansky and many others have observed, biology and evolution are conceptually inseparable – a fact that helps explain the popularity of the evolutionary definitions the pessimists single out for special attention. Evolution is not like the strong nuclear force – it is an inherently stochastic process, driven by random variation. Thus, it is literally not possible to have evolution without first having variation. If biology depends on evolution, and evolution depends on variation, we should not *expect* cleanly distinguishable biological entities. Biology, as anyone who has tried to identify organisms in the wild from a field guide understands, blurs all lines.

Another uniquely biological consideration is that natural selection exerts its influence on the basis of *functional* characteristics and does not 'see' the causal structures making these functions possible. A harmless butterfly that mimics a poisonous cousin will do well, not because of the biochemical details of its pattern formation system, but because of the functional effect such a system produces: mimicry. The functional nature of biological explanation causes two additional complications for anyone wishing to impose clean necessary and sufficient conditions on biology. First, it is inherently difficult to give a precise account of just what we mean when we talk about functions (Allen & Bekoff 1995; Wouters 2003). Second, given the enormous complexity and diversity of biological mechanisms, there will usually be many specific causal mechanisms that can produce a given function (hence evolutionary phenomena such as convergence).

The tension between universality and mechanism

"*It is the mark of an educated mind to rest satisfied with the degree of precision which the nature of the subject admits and not to seek exactness where only an approximation is possible*."
- Aristotle

Gayon (2010) observed that definitional pessimists seem to feel more comfortable with a *mechanistic* approach. The basic idea here is that to *explain* a phenomenon is to give an account of a detailed causal account (a mechanism) capable of *producing* it. While Cleland never explicitly identifies herself as mechanist, many of her claims seem to imply something of the sort[4]. For example, in several places, Cleland intimates that part of the job of a proper general theory of living systems is to specify a causal system that encompasses the biochemical details of inheritance and metabolism:

"*…some of the molecular building blocks of proteins and nucleic acids could have been different. Indeed, it is an open question as to whether all life (wherever it may be found) is constructed of proteins and nucleic acids. This question is difficult to answer outside the context of a general theory of living systems, something that we currently lack*." (Cleland & Copley 2005, p. 166)

Similarly, her insistence on microbes as counterexamples to evolutionary definitions of life is motivated largely by the fact that their detailed mechanisms of inheritance and biochemistry are different from those of multicellular organisms.

While I certainly do not want to deny that being able to describe the detailed causal mechanism behind a phenomenon is an excellent indicator that we understand its causal basis, being able to do this should not be considered a *necessary* condition for scientific explanations. One reason for this is that some questions are simply too general to allow for such answers *in principle*. Take the case of astrobiology: the pessimists rightly point out that any definition of life must be capable of encompassing the enormous diversity of living systems we are likely to encounter elsewhere in the universe. What they seem not to appreciate, however, is that there is an ineliminable tradeoff between universality and specificity. If what you seek is an account of life that can cover all life in the universe, then it

[3] The states of any digital signal will be easier to differentiate cleanly than a similar analogue one by virtue of the relatively limited number of states its components can adopt. Thus, polygenic traits tend to produce continuous variation, while single gene traits produce highly discrete phenotypes.

[4] At a recent conference, Eörs Szathmáry responded to Cleland's presentation that we should seek a *minimal* definition of life, which would of necessity be fairly abstract. Cleland replied by saying that such a description would be 'too high level' and that the real problem is 'getting molecules to do what you want.' Szathmáry ended the exchange with the Bernard Shaw quip, 'Although I cannot lay an egg, I am a very good judge of omelets.' (Smith 2015).

https://doi.org/10.1017/S1473550416000021 Published online by Cambridge University Press

must be quite general, since a strong demand for mechanical specificity can only be satisfied by imposing non-universal boundary conditions. Thus, the biochemical mechanisms of life *on Earth* are made possible only by boundary conditions such as the existence of oxygen and liquid water that may not shape life elsewhere in the universe.

This tradeoff between universality and specificity seems to be a perfectly general problem that applies to the physical sciences as well. Chemistry is the paradigm example of a science whose basic theory is mostly structural, which is part of the reason the pessimists use the emergence of chemical theory from alchemy as their inspirational story. But imagine asking a chemist to describe a 'universal chemistry' applicable to all reactions that might occur anywhere in the universe. Since she cannot assume the existence of any particular kinds of molecules, temperature regimes, ambient pressures, etc., she could only describe chemical possibilities at a very high level of abstraction. There are universally true things she can say, of course, but they will revolve around highly general, non-structural properties (e.g., basic thermodynamic considerations). The richly detailed chemical mechanisms learned in an organic chemistry class on Earth are not sufficiently robust to make the cut. In short, an account that is at once highly specific and highly general is simply not possible.

It can even be argued that part of our difficulty in finding an adequate definition of life may lie in our emphasis on causal detail over more general, functional properties. Allow me to briefly consider a variation of Putnam's (1975) famous 'twin earth' thought experiment to explore this claim. Imagine we discover an alien world that has a complex ecosystem much like ours, at least in broad functional terms. They have evolved millions of different species, both single celled and multicellular, that form excruciatingly complex ecosystems. They have predation, competition, nutrient cycling and many of the other general characteristics we see in terrestrial life. Of course, there is nothing *exactly* like a human or a dolphin or an *Escherichia coli*, though there are sometimes organisms that fulfil generally similar roles in generally similar ways. Further, let us suppose it turns out that the aliens have *fundamentally* different causal mechanisms underlying these functional similarities. Not only do they not have nucleic acids or proteins, they do not possess complex organic molecules *at all*. Instead, they utilize some truly exotic form of chemistry nobody on Earth had ever even speculated about as a potential basis for life. Now ask yourself, 'Are these entities alive?'

If you believe the details of the causal structure are *essential* to the definition of life, then you have two basic choices. First, you could say they are not alive – or perhaps that they are not alive in the same way as terrestrial organisms. Second, you could reserve judgment until we understand more about the comparative chemistry. In either case, the basic assumption is that the question of living status hinges on highly specific details about their chemistry, which seems a very odd claim. After all, the salient feature of the situation seems clearly to be that these organisms *do* all kinds of things we associate with living beings, and at a high level of complexity. The fact that their chemistry is different is a side note, though of course a scientifically interesting one. My guess is that most biologists would immediately grant that these aliens are alive, which shows that they think about life, when push truly comes to shove, more in terms of broad functional properties than mechanical details[5]. To be sure, they will also want to study these mechanisms to learn *how* these organisms are alive, but they are unlikely to consider the precise biochemical details germane to the question of *whether* they are alive.

So to the extent that we believe biological categories are inherently functional, we should be wary of too much causal detail in our definitions. Cleland clearly does not think much of this possibility, however:

> "*The notion that these [functional] characteristics provide the best candidates for essential properties of life rests upon the assumption that life is a functional as opposed to compositional or structural kind. But there is little empirical evidence to support the assumption that life is a functional kind. For all we know these pervasive functional characteristics of contemporary Earth life represent unreliable symptoms of more fundamental but as yet unknown properties of life*." (Cleland 2012, p. 130)

Perhaps life is not a functional kind, of course – that question is beyond the scope of this paper. But the hypothesis that the essence of life is structural seems no better supported empirically. Indeed, the lack of a unified account of life, despite the enormous emphasis placed on elucidating structural details in modern biology, at least suggests that we may be looking in the wrong place. Cleland herself admits this, and even hypothesizes that the key to a general definition of life may be some property unknown to modern biological science:

> "*We also can't rule out the possibility that the most important characteristics of life have yet to be discovered. The characteristics traditionally held up as essential to life may be little more than potentially unreliable symptoms of more fundamental but as yet unknown properties.*" (Cleland 2007, p. 849)

Since biology has no analogue to the hidden variables proof of quantum physics, this is certainly a possibility. In the absence of any good evidence one way or the other, it seems more likely that life is inherently functional and thus the messiness is inherent, than that there is a mysterious factor which will one day make all the messiness vanish. But time will tell.

The alchemy analogy

> "*Philosophy of science without history of science is empty; history of science without philosophy of science is blind.*"
> - Imre Lakatos

The pessimists use an analogy between alchemy and attempts to define life that has a powerful intuitive appeal. The basic problem with alchemy, according to Cleland, is that it lacked

[5] Our intuitions might be different if these creatures turned out to be something like biological *robots* created by another intelligent species, etc. But we can avoid such counterexamples (if we wish) simply by requiring that entities have *evolved* to act as they do. Note that this is another *functional,* as opposed to mechanical, requirement.

https://doi.org/10.1017/S1473550416000021 Published online by Cambridge University Press

an adequate theoretical account of its objects of study. Alchemists therefore lumped together very different kinds of compounds, believing for example that ordinary water and nitric acid were the same sort of thing:

> "*Lacking recourse to molecular theory, the alchemists, who were medieval chemists, chose solvency as the essential property of water, and as a consequence identified chemical substances, for example nitric acid, that we now know are not water as water; it is not an accident that they called nitric acid 'aqua' fortis.*" (Cleland 2007, p. 849)

The 'water' the alchemists described was thus not a natural kind existing independently of their theories, but rather a category of convenience, lumping together what we know now to be very distinct kinds of things. The idea is that the alchemists were so ignorant that they considered the two to be essentially the same, which is why they used the term 'aqua' to describe both.

This certainly sounds like a damning error. But the way we think of classification here draws on our modern conception in a way that is very unfair to the alchemists. Consider the seemingly simple question, 'Did the alchemists believe that nitric acid was the 'same thing' as water?' It is actually hard to say, precisely because most alchemists did not have a corpuscular theory of matter. Whatever they meant by 'water' was therefore not the same as our modern conception and, in particular, was not tied to the kind of detailed structural account a modern chemist relies on. The primary mechanism of classification in alchemy was functional: compounds were grouped largely on the basis of shared functional properties. For example, *aqua regia* ('noble water' – a highly corrosive mixture of nitric and hydrochloric acids) was so called because it was able to dissolve the 'noble' metals (gold and platinum).

Given this, how should we characterize the alchemists' understanding of the relation between water and aqua regia? Let us look at what they knew about aqua regia. They knew that, unlike ordinary water, it was produced in a complex process of synthesis that required many different compounds other than water. And they knew that it had highly unusual properties that ordinary water does not – namely (in this case quite literally) that it was highly corrosive (Rasmussen 2014). On the other hand, they also realized that it shared properties with water: it was liquid, permeable to light, dissolves the things water dissolves, etc. The fairest thing to say is probably that they recognized water and aqua regia were very different in many ways, but they also recognized that they were similar in some ways. That is why they differentiated them, but with related names. To imply that they thought water and aqua regia were *the same thing* in a strong sense is thus clearly unfair. Surely any alchemist worth his salt would have interceded if his patron proposed taking a bath in aqua regia, and for good reasons he could explain[6].

It is something of a theme for this paper that it is often hard to define categories cleanly, and 'alchemy' is certainly no exception. Historically, what constitutes alchemy and what (if any) core beliefs its practitioners shared is extremely complicated. If we wish to be historically accurate, we cannot separate the unscientific practices of alchemy from the science of chemistry as neatly as Cleland's analogy implies. Alchemy is probably better thought of as an early stage in the development of chemistry than as an unscientific opponent vanquished by the advance of scientific method. Newman & Principe (1998) argue that the boundary between chemistry and alchemy was 'extremely diffuse at best' and thus recommend an entirely new term (*alchymy*) for the practices of this period to ward against our tendency to oversimplify, particularly in ways that flatter modernity (see also Fors 2015).

It is not even possible to cleanly divide individual *practitioners* of alchemy into the scientifically astute and the hopelessly naïve. Consider Tycho Brahe, the 16th century astronomer and alchemist. On the one hand, he has been described as "*the first competent mind in modern astronomy to feel ardently the passion for exact empirical facts*" (Burtt 1925). He is rightly lauded for the exacting care he took with his astronomical observations, which were far more accurate than those of his contemporaries. If we picture Brahe as a man struggling to make extremely precise measurements of the positions of celestial bodies, he seems thoroughly scientific. But this very same Brahe was also an ardent astrologer and is said to have employed a household dwarf to predict the future. If we focus instead on these beliefs, he seems very far indeed from a modern scientist. So if one asks, 'Was Tycho Brahe a scientist?' the only adequate answer is the unsatisfyingly ambiguous, 'That depends on what you mean by *scientist*.' Similar things could of course be said of other scientific luminaries such as Sir Isaac Newton, which really should not be so surprising, since the boundaries of social categories, like those of biology, tend to be blurry.

Finally, there is a strand of history and philosophy of science that believes we focus far too much on the grand *theories* of science (e.g. atomic theory) when we tell our histories (Ihde 1991; Davis 1993). The fact of the matter is that much of our scientific progress owes more to the development of techniques and equipment than to theory. Certainly in this respect, the alchemists deserve enormous credit, as chemistry would have had very little data to work with if not for the discoveries of generations of alchemists who preceded them. To again take up the example of aqua regia, it was not possible to prepare highly concentrated acid solutions until the invention of the retort in the early 14th century. Alchemists had to first develop, through a long series of prototypes, the equipment to produce such solutions. Then they conducted further experiments using this equipment to create a whole range of different compounds, including aqua regia. Finally, they experimentally determined

[6] Indeed, it is tempting to push a point and defend the alchemists' terminology further by pointing out that, even with a modern conception of chemical composition, they had a point when they classified aqua regia as a kind of water. With the exception of the fuming acids, which contain no water at all and were far beyond the capabilities of the alchemists to produce, acids are typically *aqueous* solutions. Therefore, a philosopher who says, 'Nitric acid is HNO_3' is putting forward an *idealization* that does not match the ordinary usage of even a modern chemist.

https://doi.org/10.1017/S1473550416000021 Published online by Cambridge University Press

the various properties of the compounds and used these properties to classify them as best they could. Was this enough to make them scientists in the modern sense? No – but neither were they merely superstitious dabblers.

Let us set aside worries about historical accuracy at this point and assume for the purposes of argument that a clean division between the unscientific alchemists and the scientific chemists makes sense. Even then, the pessimists' joy will be short lived, because the very fact that one group is scientific and the other unscientific undercuts the analogy. Analogies only work to the extent that they compare similar things – if two things are similar in many aspects we can confirm, it seems reasonable to assume that they are probably similar in other ways we cannot confirm. On the other hand, the more different the things being compared are, the less reasonable the assumption of conformity becomes – logicians call this *the problem of analogy*.

In this case, alchemy and attempts to define life are similar in their inability to come to a clean consensus concerning the ultimate nature of the objects they are studying. The pessimists would thus have us conclude that the prospects of an adequate scientific definition of life today is no better than the prospects for an alchemical definition of matter were 400 years ago[7]. But, whatever we think of the scientific status of alchemy, the new fields in biology that are pushing to define life are *clearly* scientific in the fully modern sense. Investigators in these fields are trained in traditional scientific disciplines and pursue serious scientific investigations using the same types of methods other scientists employ. Most critically, they have explicitly disavowed the pre-scientific metaphysical views alchemists often espoused – rejecting, for example, non-natural forces as legitimate explanatory tools. So the analogy to alchemy is intuitively powerful for the same reason it is unfair: modern biologists should not be labelled 'unscientific.'

The analogy of planetary science and the $N = 1$ problem

> "*Life is made up of a series of judgments on insufficient data, and if we waited to run down all our doubts, it would flow past us.*"
> – Learned Hand

I would like to suggest a more realistic analogy for modern attempts to define life: the state of planetary science 25 years ago. Just as with the disciplines pushing to define life today, planetary science then was being done by investigators with a fully modern scientific outlook. But it also faced problems very similar to those life researchers confront today.

One basic problem biologists thinking about definitions of life have is a very limited data set. This is often described as *the $N = 1$ problem*, since all the various forms of life we know are descended from a single origin event and thus can, in some sense, be described as a single data point. Planetary science was in much the same boat until 1992, since we had not confirmed the existence of even a single planet outside our own Solar System. Of course, we had (limited) access to our neighbouring planets, and these were studied as extensively as we could. Still, there was an important sense in which the data was clearly inadequate, since any honest planetary scientist would have to admit that the data available was too limited to allow truly robust tests of theory. It was thus an open question as to how well planets in other systems would conform to the theories of planetary science. So to the extent the $N = 1$ problem exists now for biologists attempting to define life, it existed then for planetary science.

This all began to change with the confirmation of the first extra-solar planet. Since then, we have discovered some 2000 exoplanets and learned a great deal concerning planetary distributions and diversity. For example, we now know that rocky worlds like Earth are more common than we thought, habitable zones are larger than anticipated and more planets are in 'unusual' orbits (Winn & Fabrycky 2015). We have even found cause to tweak our definition of a planet, with the resulting demotion of Pluto[8]. So we certainly know much more about planets now than we did 25 years ago. On the other hand, this looks much more like incremental progress than revolution. What we have done is fill in some of the *details* underlying the general theory of planetary science, but in ways that have not really caused us to reassess the accuracy of what we knew before in any fundamental way. Nobody has questioned whether planetary science prior to 1992 deserved to be called 'science' or suggested that earlier attempts to define a planet were philosophically misguided because they lacked an adequate empirical basis. Indeed, the real story here is the extent to which planetary science, for all its difficulties, got things *right*.

There is a reason for this that is instructive: planetary science was based, not just on simple induction from observing planets near us, but on theoretical extrapolations from many other scientific disciplines. For example, a spherical shape was part of the definition of a planet not simply because we noted all the solar planets happened to be spherical, but because our knowledge of material science and physics led us to conclude that any sufficiently large body formed by orbital accretion would adopt a spherical shape under the influence of gravity. Deciding to add the stipulation that a true planet must clear its orbit of debris thus did not shake the core of the theory.

There is every reason to expect something similar as synthetic and astrobiology expand our data set of living systems. Biology is based on evolution and the general theory of evolution draws on a number of related disciplines in ways that confirm its general predictions. As a consequence, while there is little doubt we will discover a wealth of new detail concerning how living systems can work, there seems little reason to expect our basic conception of how life comes to be will be fundamentally altered. For example, there are good reasons to think that any natural system we might be tempted to describe as alive

[7] And perhaps also that the *reason* for their failures is the same: an inadequate account of the mechanical details underlying the functional properties used in classification.

[8] An excellent discussion of the difficulties surrounding scientific classification in general and the case of Pluto in particular can be found in Dick 2013.

https://doi.org/10.1017/S1473550416000021 Published online by Cambridge University Press

must share functional properties that have long held pride of place in our definitions, such as the capacity to metabolize and evolve. As long as these are described at a sufficiently high level of generality, there seems no reason to expect they will prove unrepresentative of life elsewhere. One way to put this is that the idea that life beyond Earth will be fundamentally *different* than terrestrial life should be treated like any other hypothesis for which we require evidence. The fact that the universe will likely be very diverse in its details is not good evidence that it will be as diverse in more general properties.

Like most pessimists, Cleland makes much of the $N = 1$ problem:

> "*One cannot safely generalize to all life, wherever and whenever it may be found, from a sample of one. To do so would be a bit like trying to come up with a theory of mammals based solely on observations of zebras. It is unlikely that someone faced with this task would focus on their mammary glands because they are characteristic only of the females. A more plausible candidate would be their ubiquitous stripes. Yet as biologists have discovered, the mammary glands of female zebras are far more relevant to their nature as mammals than their stripes.*" (Cleland 2013b, p. 370)

While it is probably true that we could not learn much from a single zebra, a persistent biologist might be able to learn a surprisingly amount from a *population* of zebras. The bottom line is that, the more inclusive the data set becomes, the more variation it represents, and thus the more it reveals about the evolutionary dynamics of the system which created it. So rather than calling life on Earth a single data point, it is much more accurate to call it a single data *set*. Data sets, for example a population of zebras, can exhibit significant variation of the sort needed to support scientific generalizations, though of course we never know for certain whether these can be safely universalized. When scientists had only Earth to study, their ability to extrapolate to the characteristics of planets in general was severely constrained. But by modern times, we had a data set that included dozens of planets and similar bodies with very different characteristics, greatly improving the prospects of correctly identifying universal principles.

To expand on Cleland's thought experiment, what might a biologist discover if she had not just a population of zebras, but the entire genus *Equus* to study? It is not at all clear that she would not be able to deduce how zebras came to be, at least in very broad strokes, from such a data set. Now consider the actual situation we face in modern biology – biologists have something on the order of 7–10 million species to work with (Sweetlove 2011). There may well be aspects of evolutionary dynamics not captured in such a data set, but to call the whole of terrestrial life *a sample of one* is a stretch.

Microbes as exemplars

> "*For the first half of geological time our ancestors were bacteria. Most creatures still are bacteria, and each one of our trillions of cells is a colony of bacteria.*"
>
> - Richard Dawkins

Cleland and her colleagues talk a lot about the importance of studying microbes for insights into the nature of life. It is thus worthwhile to look at three types of claims they put forward here to gain further insight into their position. First, there is the claim that, to the extent we must extrapolate from the available terrestrial evidence, we are well-advised to focus on microbes. As Dawkins notes, microbes were the only life on earth for most of its history, and even today they vastly outnumber their multicellular cousins. Thus, it seems reasonable to use microbial life as our basic template for alternate forms of life, whether in space or the laboratory. This is certainly an unobjectionable point – indeed, that is the problem. Pretty much everyone seriously engaged in thinking about the nature of life already accepts this as a given. The question is thus not *whether* a microbial focus is warranted, but rather *exactly how* such a focus impacts our conception of life.

Cleland and her collaborators argue that focusing on microbes changes everything because they introduce all kinds of 'weird' biological mechanisms like lateral gene transfer. For example, they argue that, if we really take the microbial example seriously, we cannot defend a traditional *Darwinian* account of life. Consider Cleland's discussion of Carl Woese's musings on the origin of life:

> "*These proto-cells could not evolve in a Darwinian fashion because they lacked the sophisticated genetic apparatus of modern cells, viz., the complex cooperative arrangement between proteins and nucleic acids as mediated by ribosomes…Regardless of whether one is sympathetic with the details of his account, Woese is almost certainly correct in rejecting the claim that the complex cooperative arrangement between proteins and nucleic acids—the molecular foundation for the phenotype–genotype distinction so important to Darwinian evolution by natural selection as traditionally understood—emerged full blown at the time that life originated.*" (Cleland 2013b, p. 377)

Perhaps it will come as no surprise at this point that responding to this objection involves us in more terminological fuzziness. This is unfortunately unavoidable, since what 'Darwinian' means varies greatly from one author to another. Some use it as shorthand for the broad sweep of evolutionary theory in general, others to indicate views that emphasize evolutionary features Darwin himself thought important (e.g., gradualism) and still others as shorthand for some (often unspecified) version of 'prevailing evolutionary beliefs.' Here, Cleland seems to be taking the latter path, identifying Darwinian evolution as evolutionary theory with the essential inclusion of vertical gene transmission and the specific details of gene–protein interactions in modern organisms. If we accept this version of what it means to be *Darwinian*, it makes sense to conclude that, to the extent we want our account of life to include microbes and systems of the sort Woese envisions, our account of life cannot be Darwinian.

This seems to be another instance of the pessimists' mechanistic bent, as they are identifying 'Darwinian' with very specific causal mechanisms. But one problem we quickly run into when we think of 'Darwinian' this way is that it will look as if evolutionary biology is in a constant state of crisis each time biologists debate some new wrinkle. This is a strategy often employed by creationists to discredit evolutionary theory in general and in fact Dembski (2002) recently cited Woese in exactly this way, suggesting that his ideas indicate dissent within the biological community concerning the adequacy of evolutionary theory[9]:

> "*There is a question about the extent of evolution, but that is a question being raised by non-ID scientists. Carl Woese in the Proceedings of the National Academy of Sciences just a few weeks ago published a piece where he explicitly rejects common descent*." (Dembski 2002)

Myers (2008) offers a succinct reply to Dembski on this point that works equally well with mechanically minded definitional pessimists:

> "*Woese recognizes a pre-Darwinian period, a Darwinian threshold and a Darwinian period. In the pre-Darwinian period, massive horizontal gene transfer made it impossible for individual species to evolve, however the Darwinian processes of "amplification, variation and selection" still played a role. In other words, "pre-Darwinian" referred to a period in which speciation was impossible, not a period in which the Darwinian processes did not play a role*." (Myers 2008, p. 1)

So Woese's views are not Darwinian only if 'Darwinian' is taken to refer to a highly specific version of evolutionary theory, which is not what Woese intended. And however we parse Woese's words, there is certainly no conflict between his theory and the general theory of evolution.

Precisely in the same way, there is no conflict between most of the evolutionary accounts of life that are actually being debated and what their advocates take to be 'Darwinian.' In short, the pessimists are attacking a straw man here. Consider the most widely discussed evolutionary definition[10]: *Life is a self-sustained chemical system capable of undergoing Darwinian evolution* (Joyce 1994). There is nothing to suggest that the term *Darwinian* here implies any particular *kind* of chemical system or pattern of inheritance or that applying it to microbes or Woese's proto-cells would pose a special problem. Other accounts are even more general – as when Bedau (1996) discusses life in terms of a "*supple, open-ended evolutionary process that perpetually produces novel adaptations*." At its most basic, evolutionary theory simply predicts that when systems with heritable adaptive variation are exposed to adaptive challenges, they will tend to evolve to meet those challenges. It does not really matter to the general theory *how* the inheritance system is realized, though of course different mechanisms will produce different patterns of inheritance (a specific detail only important for answering certain specific kinds of questions).

Their emphasis on the importance of microbes led Cleland and Copley to hypothesize a 'shadow biosphere' on Earth, composed of microbes descended from a separate origin event and thus potentially quite distinct from life as we know it:

> "*The discovery of a shadow Terran biosphere would have profound scientific and philosophical ramifications. It is clear that life as we know it on Earth has a common origin, which means that we are currently limited to a single sample of life. One cannot generalize on the basis of a single sample. In order to formulate a truly general theory of living systems we need examples of unfamiliar forms of life. Although we have good theoretical reasons for believing that life on Earth could have been at least modestly different in its biochemistry and molecular architecture, we do not know how different it could have been. It is important that we do not artificially constrain our thinking about the origin of life on Earth and the possibilities for extraterrestrial life on the basis of a limited and possibly very misleading example of life*. (Cleland & Copley 2005, p. 171)"

A shadow biosphere is certainly an intriguing possibility and it is hard to imagine a life researcher who would not agree that its discovery would be monumentally important. As such, this is a possibility that certainly deserves more study.

On the other hand, the possible existence of a shadow biosphere should really impinge on this debate as an argument *against* definitional pessimism. Bringing a shadow biosphere into the discussion *forces* us to answer the question of how life should be defined. A *bio*sphere is, by definition, composed of living entities. Therefore, a willingness to apply that label to some unusual system we discover presupposes that its components meet *some* definition of life, however tentative. What then is that definition? Without some notion of what we are talking about, we are free to reject any putative example of a biosphere as spurious[11]. So the possibility of a shadow biosphere is actually a powerful motive to *develop* a clear account of what life is.

Finally, Cleland often points out that focusing on 'weird' examples of life (like microbes) will help us keep an open mind and avoid the straightjacket of a priori preconceptions:

> "*In this context, one cannot help but wonder whether contemporary thought about the nature of life is being held hostage to an antiquarian Aristotelian conception of life. This might explain why biologists have been unable to identify explanatorily and predictively powerful, distinctively biological generalizations despite the extensive body of knowledge they have accumulated about familiar Earth life over the past century…Their absence may reflect little more than a commitment to an inadequate ontology; distinctively biological regularities may exist but be unrecognized as such because we are carving up the phenomenon of life in unhelpful ways…this concern is exacerbated when one reflects that, in keeping with Aristotle's approach, most attempts to generalize about life have been founded upon what biologists now recognize is a rare and exotic form of Earth life, viz., multicellular plants and animals*." (Cleland 2013b, p. 370)

[9] I do not mean to imply that the definitional pessimists are trying to give comfort to the creationists, but whenever one identifies a general theory too closely with its details, this kind of objection is common.

[10] This is often referred to as 'the NASA definition' despite never having been adopted by NASA.

[11] No doubt any initial claim to have discovered a shadow biosphere will meet with precisely this problem, as skeptics will claim that all that has been discovered is some odd kind of chemical system, etc.

https://doi.org/10.1017/S1473550416000021 Published online by Cambridge University Press

> "*Because they fix (albeit tentatively) what sorts of things qualify as life and what sorts of things do not, they are just as likely to seriously mislead us as traditional definitions of 'life' if (as I shall argue is highly likely) our current scientific concept of life is unreliable. If one utilizes such a definition to guide attempts to "create" life in the laboratory or search for life on other worlds one is likely to produce or find only what one is looking for*." (Cleland 2012, p.128)

The idea seems to be that, if we agree on a particular definition of life before we really understand life, it will stifle future research.

This seems problematic for a number of reasons. First, setting aside the question of whether our philosophical ideals should conform to biology or vice versa, there is no *evidence* provided that premature definitions are somehow blocking the discovery of biological laws. Even if it is true that biology is being impeded by the lack of an adequate life definition, it seems puzzling to use this fact to conclude that we should cease all attempts to rectify the situation by seeking better definitions. Second, there are few life researchers who would argue against the desirability of looking at weird examples of life – hence the explosion of research on extremophiles. Third, an important element of the pessimists' basic position is that there are *too many* concepts of life being entertained, so how can a debate that is awash in alternative positions also be in imminent danger of premature conformity?

Pragmatic considerations

> "*Ageing's alright – better than the alternative, which is not being here*."
>
> - George H. W. Bush

It is never entirely clear what the pessimists would have us do in the absence of a general theory of biology, and it is hard to evaluate the desirability of attempting to define life without a clear view of the alternative. Yet they seem to imply that we should refuse the temptation to specify our musings about life in any formal way until such a theory materializes. If this truly is their recommendation, it is both unrealistic and unhelpful.

We can refuse to discuss what we mean by 'life', but we should be under no illusion that doing so will solve the basic problem. As Cleland & Chyba (2002) themselves allow, scientists in fields where questions about the nature of life is central *must* develop and employ concepts of life, because they cannot do their job without them. Consider the tests for extraterrestrial life performed by the Viking mission to Mars: the scientists who designed the labelled release experiment were clearly taking carbon cycling to be essential to life. Thus, when their test seemed to confirm cycling, it was taken as evidence for life on Mars. Subsequent experiments looked instead for the presence of complex organic molecules and, when these could not be found, this was interpreted as evidence against Martian life. The debate about the proper interpretation of these results continues to this day and is, to a large extent, a debate over competing concepts of life (Klein 1978).

If scientists *must* employ concepts of life, what is the advantage of forcing them to keep them implicit and unstated? If anything, preventing scientists from elaborating their ideas will *exacerbate* the difficulties, since ideas that are never made explicit cannot be subjected to the kind of rigorous critical scrutiny so central to scientific practice. Vague concepts in a scientist's head *will* continue to influence her work, but in ways *she herself* may not fully understand. Allowing our ideas to remain unspoken will also strengthen the power of the familiar, *increasing* the sorts of biases (multicellular, terrestrial) that worry the pessimists. Why should a scientist bother to look for weird chemistries that are much harder to anticipate unless she has an explicit reason to do so – like a broad definition of life? When we are explicit about our definition of life, we embroil ourselves in debate, but it is debate of the sort that can reveal and overcome bias. Vagueness can be a powerful rhetorical tool and has been used with great effect by the creationists (Smith 2001), but it is rarely in the interests of truth and certainly does not accord with the best practices of science.

On the other hand, scientists sometimes create a false precision as a way to make annoyingly complex questions empirically tractable. Thus, scientists looking to operationalize life concepts often focus on what they can easily measure and complain that (for example), evolutionary definitions of life will make the search for extraterrestrial life much more complicated. However much sympathy we have for the operational difficulties, it is critical that we keep the proper logical sequence in mind: we must begin with a clear definition of life and develop our operational procedures from this rather than the other way around. If we pick what to look for based on what is easy rather than what is informative, we are no better than the drunk looking for his lost keys under the lamp post, not because he has any good reason to think they are there, but because that it is where the light is. It is fine to demand that, any definition produce consequences that are testable *in principle*, but it is quite another to reject any definition which is merely *difficult to test at present*. Insisting on such stringent conditions would have ruled out a number of critical theoretical developments in science before they could be confirmed. Consider the predictions of a cosmic background radiation, initially theorized as early as the turn of the century, but not confirmed for 70 years – and then by lucky accident.

Another reason the pessimists offer for not forging ahead on the definitional project is the sad state of the evidence. They feel that, since all we have is a single data point, a definition of life must wait until we have much better data. Of course, if our goal is a truly *ultimate* definition, where we are supremely confident we have captured the essence of the phenomenon, this is clearly true. But on a more pragmatic level, it is instructive to think about what waiting for the data to emerge would mean *in practice*. At a minimum, a good data set would have to span different origins of life on different planets in different planetary systems. Only then would we have the empirical grounds to be truly confident extending our principles beyond the frozen accident of Earth's peculiar history.

But this is far easier said than done. In astrobiology, even the simplest investigations will occupy scientists for decades. For example, both Enceladus and Europa have recently been found to vent water vapour into space, presenting about as ideal a situation for the collection of biologically promising samples as could be wished. But we will still have to build probes to make the long journeys, collect the samples and subject them to initial analysis. We will likely have to return the samples to Earth for definitive answers, since (as the Viking mission showed) the experiments possible on an interstellar probe are very limited (McKay 1997). So it could easily take 25–50 years before we have really good evidence for or against life within our own Solar System.

And this is just the beginning. Alien life on Mars or Europa, even if it represents an independent evolutionary origin, will only take us so far. We are ultimately going to need to test for life in other systems, which will require missions spanning hundreds, if not thousands, of years with any technology we can currently envision. We will have to design experiments that can survive such incredibly long journeys through the cold of space, and then operate flawlessly on planets in conditions the designers were largely ignorant of. Returning samples to Earth for more analysis will double the time needed to complete the process. Finally, we will have to repeat this procedure as many times as necessary to acquire a reasonable data set spanning the potential diversity of alien life, which could easily require 20 or more missions. It would be wonderful if this was not the case, of course, but it is. We are thus forced to ask a very pragmatic question: *Is it reasonable to expect science to wait 1000 years or more to start work on a definition that is indispensable to the cutting edge of biology right now?*

Concluding remarks

"*We must just KPO.*" – Churchill

The definitional pessimists as represented by Cleland and her colleagues advance a number of arguments. They espouse a traditional notion of *definition* in which no definition is said to be adequate if it cannot identify necessary and sufficient conditions that produce neat and unambiguous categories. To this they add an additional requirement that definitions must specify *detailed causal mechanisms* explaining how the features of these categories are generated. They also argue for a more inclusive approach to the exemplars we use, particularly in our search for life beyond Earth, and particularly with respect to the status accorded our microbial cousins. Finally, they warn that, until a truly general theory of biology emerges, attempts to define life will be as confused and pointless as the alchemists' attempts to determine the nature of matter in the absence of atomic theory.

The definitional pessimists are driven by the need to formulate a definition of life that gets it right. In that, we are entirely of one mind. There is nothing wrong with postulating a hidden uniformity behind apparent variation, and in fact *anyone* who advocates a definition of life is attempting to do precisely this. But a recurring theme of this paper is just how *messy* definition can be, whether we are talking about life, Darwinism, or alchemy. The claim that one approach to scientific definition is superior to all others is hardly uncontroversial among philosophers. And even if such definitions were *the* appropriate ideal in other sciences, there are good reasons to suspect that forcing them on *biological* entities, produced by evolutionary forces that thrive on variation, will do more to confuse than enlighten.

There is also nothing wrong with emphasizing that a mechanism explaining in great detail how a particular system works is excellent evidence that we truly *understand* that system. It is also true that biologists often do not accord microbes and their unique mechanisms the prominence they deserve. But neither claim is especially controversial as there are a great many people pursuing definitions of life who would not object to either. The matter is quite different however, if we go further and say either that detailed mechanisms are *required* for *any* adequate definition or that it is impossible to create 'Darwinian' accounts that do justice to microbial biology. The former represents an impossible demand, since the generality we require of *any* truly universal account is inconsistent with a high level of causal detail. The latter is only true if we adopt a highly restricted concept of Darwinism that is not representative of most defenders of evolutionary accounts of life.

Finally, the pessimists invite us to consider that those who attempt to define life are in much the same position as the medieval alchemists: not only are their attempts doomed to fail without a general theory of biology, but such attempts may *impede* progress by leading us down false paths and encouraging group think. Yet since the development of modern chemistry clearly owes much to the alchemists' efforts, perhaps the analogy is not as unwelcome as it initially appears. To the extent the comparison stings, it does so by suggesting that all such attempts are inherently *unscientific*, which is unfair to *all* the disciplines concerned. In any event, a much better analogy can be found in the recent development of planetary science, which overcame its own lack of data to vindicate the general aspects of its theory.

There is no getting around the fact that biology is messy, especially when it comes to questions about the nature of life. The difference between definitional pessimists and optimists is really about how we should *deal* with this messiness. Pessimists would give up on the definitional enterprise until the empirical picture is clearer, while optimists advocate pressing on despite the difficulties. In addition to defusing many of the arguments the pessimists offer, I suggest two reasons for maintaining a sense of optimism, at least as a heuristic. First, if we accept that scientists engaged in cutting edge disciplines like synthetic biology and astrobiology need *some* concept of life to do their work, then the pessimists' recommendation is really that these scientists should use implicit rather than explicit concepts. It is hard to see, how encouraging people to keep their working concepts private and vague will foster progress – indeed, there is every reason to expect as it will only make many of the problems the pessimists identity *worse* by circumventing the process of public critique that is central to science. Second, if we are to wait until we have an unambiguously adequate empirical basis for a universal account of life, we will have to be very patient indeed. The pessimists complain that some definitions of life are problematic because they are difficult to operationalize, but fail to consider how much worse the situation is if we are forced to wait literally *thousands* of years to even begin discussions about what it is that we are studying right now.

Science is a messy business, despite the best efforts of philosophers to tidy it up. But it is also an exercise in epistemological optimism and scientists are necessarily the sorts of people who doggedly pursue explanations even when there seems little immediate hope of success. It is thus absolutely central to the spirit of science that we resist giving in to pessimism when nature refuses to give up her secrets as quickly as we would like. New ways of thinking tend to come about only when we face such challenges head on. Therefore, however messy and confused things may be, science really has no choice but to follow Churchill's advice and just 'keep plodding on.' In fact, we would be well advised to produce *more* definitions of life and wholeheartedly embrace the debate that ensues:

> "*Of course, the discovery of some system or phenomenon that we unanimously agree to regard as an alternative form of life, when it comes (if it eventually comes), will have deep implications and change profoundly our worldviews and our conception of the living. But, meanwhile, efforts to put together explicitly, in a distilled way, what our day-to-day increasing biological knowledge is telling us about the actual concept of life should not be abandoned. Quite the contrary, they should be encouraged, as part of what biology is really lacking right now: more encompassing approaches that contribute to integrate the huge amounts of data and relevant information being continuously generated.*" (Ruiz-Mirazo *et al*. 2010, p. 206)

Acknowledgements

This paper has benefitted from discussions with a number of researchers, but special thanks are owed to John Baross, Steven Dick, Carlos Mariscal and Lucas Mix. An anonymous reviewer also helped sharpen the focus of the paper in several important ways, for which I am grateful.

References

Allen, C. & Bekoff, M. (1995). Biological function, adaptation, and natural design. *Philos. Sci*. **62**(4), 609–622.

Bedau, M. (1996). The nature of life. In *The Philosophy of Artificial Life*, ed. Boden, M., pp. 332–356. Oxford University Press, London.

Bedau, M. & Cleland, C. (2010). *The Nature of Life: Classic and Contemporary Perspectives from Philosophy and Science*. Cambridge University Press, London.

Boyd, R. (1991). Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philos. Stud*. **61**(1/2), 127–148.

Burtt, E.A. (1925). *The Metaphysical Foundations of Modern Physical Science*. Trench, Trubner & Company, New York.

Cartwright, N. (1980). The truth doesn't explain much. *Amer. Phil. Q*. **17**(2), 159–163.

Chyba, C.F. & Mcdonald, G.D. (1995). The origins of life in the solar system. *Ann. Rev. Earth Planet. Sci*. **23**, 215–249.

Cleland, C. (forthcoming). *The Quest for a Universal Theory of Life*. Cambridge University Press, London.

Cleland, C. (2004). Why it is a mistake to define 'life'. In *Artificial Life IX Workshop Proc.*, ed. Pollack, J., Bedau, M., Husbands, P., Ikegami, T., & Watson, R.A., pp. 93–95. MIT Press, Cambridge, MA.

Cleland, C. (2006). Understanding the nature of life. In *Life as We Know It*, ed. Seckbach, J., pp. 589–600. Springer, Dordrecht.

Cleland, C. (2007). Epistemological issues in the study of microbial life. *Stud. Hist. Phil. Biol. Biomed. Sci*. **38**, 847–861.

Cleland, C. (2012). Life without definitions. *Synthese* **185**, 125–144.

Cleland, C. (2013a). Is a general theory of life possible? *Biol. Philos*. **28**, 189–204.

Cleland, C. (2013b). Is a general theory of life possible? *Biol. Theory* **7**, 368–379.

Cleland, C. (2013c). Conceptual challenges for contemporary theories of the origin of life. *Curr. Org. Chem*. **17**, 1704–1709.

Cleland, C. & Chyba, C. (2002). Defining 'life'. *Orig. Life Evol. Biosph*. **32**(4), 387–393.

Cleland, C. & Chyba, C. (2007). Does life have a definition? In *Planets and Life: The Emerging Science of Astrobiology*, ed. Sullivan, W.T. & Baross, J., pp. 119–131. Cambridge University Press, London.

Cleland, C. & Chyba, C. (2010). Does life have a definition? In *The Nature of Life*, ed. Bedau, M. & Cleland, C., pp. 326–339. Cambridge University Press, London.

Cleland, C. & Copley, S.D. (2005). The possibility of alternative microbial life on earth. *Int. J. Astrobiol*. **4**, 165–173.

Cleland, C. & Zerella, M. (2013). What is life? In *The Philosophy of Biology*, ed. Kampourakis, K., pp. 31–48. Springer, Dordrecht.

Cooper, G. (1996). Theoretical modeling and biological laws. *Philos. Sci*. **63**, Supplement, S28–S35.

Davies, P., Benner, S., Cleland, C., Lineweaver, C. & McKay, C. (2009). Signatures of a shadow biosphere. *Astrobiology* **9**, 241–249.

Davis, B. (1993). Analytic chemistry and the 'big' scientific instrumentation revolution. *Ann. Sci*. **50**(3), 267–290.

Dembski, W.A. (2002). Eugenie Scott and the NCSE. *Theism.net*. http://www.theism.net/article/26 (accessed 18 October 2015).

Dick, S.J. (2013). *Discovery and Classification in Astronomy*. Cambridge University Press, London.

Dieckmann, U. & Doebeli, M. (1999). On the origin of species by sympatric speciation. *Nature* **400**(6742), 354–357.

Donoghue, M.J. (1985). A critique of the biological species concept and recommendations for a phylogenetic alternative. *Bryologist* **85**, 172–181.

Dyson, F. (1985). *Origins of Life*. Cambridge University Press, London.

Elgin, M. (2006). There may be strict empirical laws in biology after all. *Biol. Philos*. **21**(1), 119–134.

Fleischaker, G.R. (1990). Origins of life: an operational definition. *Orig. Life Evol. Biosph*. **20**, 127–137.

Fors, H. (2015). *The Limits of Matter*. Chicago University Press, Chicago.

Fry, I. (2000). *The Emergence of Life on Earth*. Rutgers University Press, New Brunswick.

Gayon, J. (2010). Defining life. *Orig. Life Evol. Biosph*. **40**(2), 231–244.

Ghiselin, M.T. (1987). Species concepts, individuality and objectivity. *Biol. Philos*. **2**(2), 127–143.

Hacking, I. (1991). A tradition of natural kinds. *Philos. Stud*. **61**(1–2), 109–126.

Hempel, C. (1966). *Philosophy of Natural Science*. Prentice-Hall, Englewood Cliffs.

Ihde, D. (1991). *Instrumental Realism*. Indiana University Press, Bloomington.

Joyce, J. (1994). Forward. In *Origins of Life: The Central Concepts*, ed. Deamer, D.W. & Fleischaker, G.R., pp. xi–xii. Jones and Bartlett, Burlington.

Keosian, J. (1974). Life's beginnings – origin or evolution? *Orig. Life* **5**, 285–293.

Klein, H.P. (1978). The viking biological experiments on mars. *Icarus* **34**(3), 666–674.

Machery, E. (2012). Why I stopped worrying about the definition of life. *Synthese* **185**(1), 145–164.

Mayden, R.L. (1997). A hierarchy of species concepts. In *Species: The Units of Diversity*, ed. Claridge, M.F., Dawah, H.A. & Wilson, M.R., pp. 381–423. Chapman and Hall, London.

Mayr, E. (2000). The biological species concept. In *Species Concepts and Phylogenetic Theory: a Debate*, ed. Wheeler, Q.D. & Meier, R., pp. 17–29. Columbia University Press, New York.

Mayr, E. (2007). *What Makes Biology Unique?* Cambridge University Press, London.

McKay, C. (1997). The search for life on Mars. *Orig. Evol. Biosph*. **27**, 263–289.

Mix, L. (2015). Defending definitions of life. *Astrobiology* **15**(1), 15–19.

Myers, P.Z. (2008). Woese, the Darwinian Threshold and Intelligent Design. *Panda's Thumb*. http://www.pandasthumb.org/archives/2008/03/woese-the-darwi.html (accessed 18 October 2015).

Newman, W.R. & Principe, L.M. (1998). Alchymy versus chemistry. *Early Sci. Med*. **3**(1), 32–65.

Noor, M.A.F. (2002). Is the biological species concept showing its age? *Trends Ecol. Evol*. **17**(4), 153–154.

Pirie, N.W. (1937). The meaningless of the terms life and living. In *Perspectives in Biochemistry*, ed. Needham, J. & Green, D., pp. 11–22. Cambridge University Press, London.

Putnam, H. (1975). *Philosophical Papers*, Vol. **II**. Cambridge University Press, London.

Rasmussen, S.C. (2014). *The Quest for Aqua Vitae*. Springer, New York.

Rosenberg, A. & McShea, D. (2007). *Philosophy of Biology*. Routledge, New York.

Ruiz-Mirazo, K., Peretó, J. & Moreno, A. (2010). Defining life or bringing biology to life. *Orig. Life Evol. Biosph*. **40**, 203–213.

Smith, K.C. (2001). Appealing to ignorance behind the cloak of ambiguity. In *Intelligent Design Creationism and its Critics*, ed. Pennock, R., pp. 705–735. MIT Press, Cambridge.

Smith, K.C. (2015). From notes taken at the Second Conf. on the History and Philosophy of Astrobiology in Hoor, Sweden, 7 May.

Sweetlove, L. (2011). Number of Species on Earth Tagged at 8.7 Million. *Nature Online*, http://www.nature.com/news/2011/110823/full/news.2011.498.html (accessed 18 October 2015).

Via, S. (2001). Sympatric speciation in animals. *Trends Ecol. Evol*. **16**(7), 381–390.

Winn, J.N. & Fabrycky, D.C. (2015). The occurrence and architecture of exoplanetary systems. *Ann. Rev. Astron. Astroph*. **53**, 409–447.

Wouters, A.G. (2003). Four notions of biological function. *Stud. Hist. Philos. Sci*. **34**(4), 633–668.