DEFINING 'DEAD' IN TERMS OF 'DIES'

[Draft of September 2005. Please do not cite or circulate without permission.]

Cody Gilmore

Table of Contents

1. **Introduction**
2. **Precursors to Rosenberg's and Feldman's Definitions: Three Naïve Proposals**
   - **2.1 Being Dead As Requiring Nonexistence**
   - **2.2 Being Dead As Requiring Nonpresence**
   - **2.3 Being Dead As Having Ceased To Be Alive**
     - **2.3.1 The Problem of Suspended Animation**
     - **2.3.2 The Problem of Deathless Fission**
3. **A Unified Solution to the Problems of Suspended Animation and Deathless Fission: Rosenberg's Definition**
4. **Feldman's Objection to Rosenberg's Definition**
5. **Feldman's Definition**
6. **An Objection to Feldman's Definition: Things In Suspended Animation That Were Dead More Recently Than They Were Alive**
7. **A New Definition of 'Dead': Being Dead As Requiring Possession Of A Toxic Property**
   - **7.1 A First Attempt To Define Toxicity**
   - **7.2 Defining Toxicity via Necessity**
   - **7.3 Defining Toxicity via Necessity and Temporal Locality**
   - **7.4 Allowing For Non-Fatal Toxic Flickers**
   - **7.5 Denying Toxicity To "Super-Cold" Suspended Animation States**
   - **7.6 Being Dead As Requiring Possession Of A Toxic Property or Nonpresence**
   - **7.7 Being Dead As Requiring A *Death* As One's Most Recent Exit From Potential Life**
8. **Conclusion**

1. Introduction

What is it for a thing to be *dead* at a time? In this paper I attempt to formulate a satisfactory definition of the concept 'x is dead at time t'. I begin by considering three "naïve" proposals – viz.:

PD1 x is dead at t =df. (i) x is alive at some time earlier than t, and (ii) at t, x does not exist.
PD2 x is dead at t =df. (i) x is alive at some time earlier than t, and (ii) x is not present at t.
PD3 x is dead at t =df. (i) x is alive at some time earlier than t, and (ii) x is not alive at t.

I argue that these are unacceptable. I then turn to a pair of definitions that appeal to the concept of *dying* at a time ('x dies at t'). These have been proposed by Jay Rosenberg and Fred Feldman, respectively:[1]

D1 x is dead at t =df. x dies at some time earlier than t.
D2* x is dead at t =df. there is some time t* such that: (i) t* is earlier than t, (ii) x dies at t*, and (iii) x is not alive at t or at any time between t and t*.

One important virtue of these proposals, which neither Feldman nor Rosenberg explicitly mentions, is that they avoid the problems facing PD1 – PD3. However, I argue that both D1 and D2 should be rejected. (In the case of D1, I simply endorse the objection that Feldman raises.) I then examine a series of new proposals that culminates in what I take to be the best available definition of 'dead' in terms of 'dies'.

As to the question of whether my final proposal is ultimately *correct*, I take no stand. What I do claim is that my proposal is the best we can do if we take 'dies' to be the conceptually prior notion and use it to define 'dead.' Thus I want to leave open the possibility that 'dead' is actually the prior notion and can be used to define 'dies.' It seems to me that the most reasonable way to decide which notion is prior is to (i) develop the best system of definitions in which 'dies' is the prior notion, (ii) develop the best system in which 'dead' is the prior notion, and (iii) compare the resulting systems. The present paper, therefore, is intended merely as a first step – albeit an intrinsically interesting one – in a larger project.

2. Precursors to Rosenberg's and Feldman's Definitions: Three Naïve Proposals

It is often said that dead things are somehow "gone". Perhaps the underlying thought is that when a thing becomes dead, it simply ceases to exist. According to this suggestion, one *necessary* condition for being dead at a given time is failing to exist at that time. Of course, no one should

[1] Rosenberg seems to endorse D1 but never formulates his proposal in just these terms. The above formulation of D1 is due to Feldman, who considers it only to reject it. And while I take Feldman's own proposed definition to be equivalent to D2*, Feldman formulates it somewhat differently. See section 5 below.

be tempted to say that this condition is *sufficient* for being dead at a time: Pegasus fails to exist now, but it would certainly be wrong to say that the Pegasus is now dead. In order for a thing to be dead at a time, it must have been *alive at* – and hence must have *existed at* – some earlier time. This principle disqualifies the following three groups of things from counting as dead at the present time: (i) permanently non-existent objects such as the Pegasus, (ii) existent but permanently *non-living* things such as the Hope Diamond, and (iii) things, such as my first grandchild, that *will be* alive but are not alive now and were not alive at any earlier time.

2.1 Being Dead As Requiring Nonexistence

With these observations in mind, we might find it natural to suggest the following definition of 'dead':

> PD1 x is dead at t =df. (i) x is alive at some time earlier than t, but (ii) at t, x does not exist (i.e., at t, there does not exist any such entity as x).

What does it mean to say that *at t, there does not exist any such entity as x*? As I shall understand this expression, it can be paraphrased as follows:

> at t, there does not exist any such entity as x =df. at time t, the proposition [~∃y(y=x)] is true.

According to PD1, then, to say that a thing x is dead at a time t is to say that (i) the thing was alive at some earlier time, and (ii) if, at t, a speaker were to assert that there is no such thing as x, his assertion would be true.

One problem for PD1 is that it forces us to take a stand on a metaphysical controversy about the existence of past entities. To see this, consider the following view:

> Eternalism The past, present, and future and their contents are all equally real; they are all equally in existence.

Eternalists endorse the following analogy between time and space: just as there exist many things (e.g., polar bears) that don't have *this place* among their spatial locations, there also exist many

things (e.g., Plato and my first grandchild) that don't have *this time* among their temporal locations. According to eternalism, just as polar bears are no less real for being spatially distant, Plato is no less real for being temporally distant. Eternalism contrasts with

> Presentism — Only the present and its contents are real; i.e., the only things that *exist* are things that are present now (intuitively, things whose careers intersect with the current time).

According to presentism, some polar bears exist but Plato and my first grandchild don't.

If eternalism is true, then a problem for PD1 arises immediately. Given eternalism, it is now true that Plato exists. But if so, PD1 forces us to say that Plato is *not* dead at the current time. And that, of course, is false: Plato *is* dead at the current time. To be sure, the case of Plato ceases to pose a problem for PD1 if, as presentists claim, there are no past entities. For in that case, it is now true that Plato does not exist. And if *that* is now true, then PD1 yields the desired result – viz., that Plato is dead at the current time.

Regardless of where we stand on the presentism v. eternalism dispute, however, we should all be able to agree that it is undesirable for an analysis of 'dead' to *entail* that eternalism is false. Other things being equal, an analysis of 'dead' should remain neutral on this dispute. But PD1 does not. So I think it should be rejected.

2.2 Being Dead As Requiring Nonpresence

One natural successor to PD1 is

> PD2 — x is dead at t =df. (i) x is alive at some time earlier than t, and (ii) x is not present at t.

Let me say a few words about the notion 'x is present at t' to which PD2 appeals. Intuitively, to say that a thing x is present at a time t is to say that the event that is x's complete *career* or

history intersects with time t.[2] Thus, if t400BC is a time in the year 400 B.C. and t2005 is the current time, then Plato is present at t400BC but not at t2005. (Plato's career intersects with t400BC but not with t2005.)

According to PD2, to be dead at a time is to fail to be present at that time but to have been alive – and hence to have been present at – at some earlier time. It should be clear that PD2 avoids the problem that led us to reject PD1. Unlike PD1, PD2 counts Plato as being dead at the current time *whether or not eternalism is true*. If past entities do not exist and eternalism is false, then of course there is no such thing as Plato, and *a fortiori* he is not present at the current time. In that case, PD2 yields the correct result: Plato is dead at the current time.

Now suppose that eternalism is true. Then it is now true that there is such a thing as Plato. Even so, it's false that Plato is *present* at the current time: Plato's career came to an end long ago and does not intersect with the current time. He and his career, though fully real and fully in existence, are entirely confined to the past. So, even if eternalism is true, PD2 yields the result that Plato is dead at the current time. In this respect PD2 constitutes an improvement over PD1.

PD2 has unwelcome commitments of its own, however. Just as PD1 forced us to take a stand on the presentism v. eternalism dispute, PD2 forces us to take a stand on the debate over:

The Termination Thesis[3] When a thing dies, it ceases to be present.[4]

---

[2] This is intended merely as a helpful heuristic. I do not mean to suggest that whenever we say that something is present at a time, we commit ourselves to the existence of events.

[3] It is more common to formulate the Termination Thesis as follows: when a thing dies, it ceases to *exist*. So formulated, it seems ambiguous: what do Terminators mean when they say, of a thing x, that it "ceases to exist" at a time t? Do they mean that, at time t, it begins to be the case that there is no such thing as x? If that's what they mean, then all eternalists ought to reject the Termination Thesis. After all, eternalists insist that *there really is such a thing as Plato*, although they agree that he's entirely confined to the past. Normally, Terminators seem to speak as if their Thesis is meant to be acceptable both to presentists and to eternalists. So presumably when they say that x "ceases to exist" at time t, they are using the phrase 'exists at' as a synonym for 'is present at' (in the manner of Sider 2001, who is very clear and unambiguous about this); what they mean is that x ceases to be *present* at t. So I suggest we frame the debate in such a way as to avoid this ambiguity.

[4] More precisely: necessarily, for any x and instant t, if x dies at t, then there is some interval T immediately following t such that x is not present at any instant in T. This allows for the possibility that a thing could die, thus ceasing to be present, and then later be revitalized, at which time the thing would begin to be present again.

A *Terminator* is someone who accepts the Termination Thesis.[5] According to Terminators, when a thing dies, what typically happens is this. There are some particles (or some matter[6]) that compose the thing in the moments leading up to its death. At the moment of death, these particles continue to be present, and they retain approximately the same macroscopic arrangement that they had before the moment of death; but they cease to compose the thing in question, and no other particles begin to compose the thing. Instead, the thing's career simply comes to an end. Thus, the living thing is present throughout the interval leading up to the moment of death, but the thing is *not* present at any moment after its death.

It bears emphasis that the Termination Thesis (TT), as I will understand it, applies to *anything* that dies. To see the force of this point, suppose that I am numerically distinct from my body, despite the fact that my body and I are currently spatially co-located and composed of just the same particles. Suppose, moreover, that my body and I will both die at time t. Then TT entails that *both* my body *and* I will cease to be present at t.

Consider the death of V. I. Lenin from the perspective of the Terminator. Call the moment at which Lenin died *time t*. Now suppose that that there were some particles – call them the *L-particles* – that composed Lenin throughout some short interval leading up to t. Then the Terminator will say the following: at t, the L-particles ceased to compose Lenin, and no other particles began to compose him. Instead, his career just came to an end, and he was not present at times after t.

---

[5] Prominent Terminators include van Inwagen 1990, Rosenberg 1998, Yourgrau 1987 and 2000, Olson 2004, and Hershenov 2005. Opponents of the Termination Thesis include Thomson 1997, Mackie 1999, Carter 1999, and especially Feldman 2000. Moreover, it seems that one must oppose the Termination Thesis if one believes (with orthodox "perdurantists" or "four-dimensionalists" such as Quine (1960: 171), Lewis (1983: 55-77), and Heller 1990) that any matter-filled region of spacetime, no matter how scattered or gerrymandered, exactly corresponds to the complete career of at least one material object. (For more on four-dimensionalism and related theses, see Sider 2001.) For consider an object whose career can be specified as follows: it consists of the second half of Plato's life and of the career of the Moon thereafter. It would seem that any object that has such a career must be alive throughout its "Platonic" phase and must die when Plato does. But it does not then cease to be present; instead it jumps, instantaneously and discontinuously, into orbit around the Earth.

[6] Henceforth I shall omit this qualification and assume that (i) all living things and all other composite material objects are composed of mereologically simple fundamental particles and that (ii) talk of matter or stuff can always be analyzed in terms of talk about things.

It would seem that parallel remarks will apply to any *other thing* that the L-particles may have composed in the moments leading up to t. If, for example, the L-particles also composed some entity distinct from Lenin such as *Lenin's body* throughout the given interval, and if Lenin's body died when Lenin did, then TT entails that Lenin's body also ceased to be present at t.

After t, most of the L-particles continue to be present and continue to be arranged in a manner that is at least superficially "Lenin-like." However, they do not compose Lenin. Nor do they compose anything – such as a certain 54-year-old *body* – that they composed in the moments leading up to t. Rather, there are two options. The first is that the L-particles fail to compose anything at all in any interval immediately following t. The second option is that the L-particles *do* compose something in some interval immediately following t. They might compose a *new* entity that suddenly "popped into existence" (began to be present) when Lenin died – e.g., a fresh corpse. On the other hand, they might compose an older entity that somehow came to be composed of the L-particles just after t despite not having died at t. (I will leave it to the reader to imagine how this might have come to happen.)

Now let us return to our discussion of PD2. If the Termination Thesis is true, then the case of Lenin poses no problem for PD2. For in that case, Lenin ceased to be present in 1924, and both he and his career are now entirely confined to the past (if they exist at all). Combined with PD2, this gives us the result that Lenin is dead at the current time, which of course is true.

Suppose, however, that the Termination Thesis is false. In particular, suppose that there are at least *some* things that continue to be present for a while after they die. It may be that *we* (e.g., you and I, Lenin and Stalin, etc.) are such things. In that case, when V. I. Lenin died, he lost the capacity for life and consciousness, began to rot, and perhaps even ceased to be a *person*, but *he did not cease to be present*. Rather, he was embalmed and put on display in Moscow, where he can be seen to this day. Indeed, this seems to be the common sense view of the matter. But if this "anti-Terminator" position is correct, and Lenin really is present at the current time, then PD2 would force us to say, absurdly, that Lenin is not dead now.

Moreover, even if each of us is by his very nature necessarily such as to cease to be present when he dies, nevertheless there might be *some* things – such as human *bodies* or trees – that live, die, and continue to be present for some time after they die. Thus, e.g., even if Lenin himself ceased to be present when he died, Lenin might have been spatially coincident with something (e.g., his body) that was alive before t, died at t, and continued to be present thereafter. But if this is correct, and Lenin's once-living body is present at the current time, then PD2 would again force us to accept an absurd conclusion. In this case it would force us to accept the absurd conclusion that Lenin's once-living body is *not* dead now.

It seems to me that 'dead' should be defined in such a way as to allow for the *possibility* of things that remain present at times at which they are dead. Just as an analysis of 'dead' should not force us to accept presentism, such an analysis should not force us to accept the Termination Thesis. Since PD2 does exactly that, it should be rejected.

2.3 Being Dead As Having Ceased To Be Alive

The following analysis of 'x is dead at t' might initially seem to be an attractive replacement for PD2:

PD3 x is dead at t =df. (i) x is alive at some time earlier than t, but (ii) x is not alive at t.

According to PD3, to be dead at a time is to fail to be alive at that time but to have been alive at an earlier time. Unlike its predecessors, this analysis correctly counts both Plato and Lenin as being dead at the present time, even if eternalism is true and the Termination Thesis is false. After all, regardless of whether Plato exists, and regardless of whether Lenin is present now, it is clear that neither of these things is *alive* now. According to PD3, therefore, they are both dead at the present time.

It seems to me that even PD3 must ultimately be abandoned, however. There are two main reasons for this: (i) the problem of suspended animation, and (ii) the problem of deathless fission.

2.3.1 The Problem of Suspended Animation

The first problem for PD3 arises from suspended animation. In laboratories, living cells and simple organisms are often frozen, preserved at very low temperatures for some time, and later thawed and thereby revitalized. As is well known, this process not merely possible: it has actually been performed on human sperm cells, ova, and blastulas, not to mention adult insects. It is uncontroversial that such things are alive both before they are frozen and after they are thawed. But what is the status of these things *while they are frozen*?

Following Michael Wreen (1987) and Feldman, I find it plausible to say that these entities are neither alive nor dead at such times. Here is Feldman:

> Organisms that enter suspended animation cease to live, but do not then die . . . We must define death in such a way that it turns out that organisms that enter suspended animation don't die. (1992a: 170)
>
> As I see it, suspended animation is a state that excludes both life and death. (1992a: 171)

The claim that organisms in suspended animation are *not alive* is supported by the fact that all of their life-functions ("metabolism . . . motion, energy transfers, syntropic increase in orderliness, etc."[7]) have been halted. The claim that such organisms are *not dead*, on the other hand, is supported by the fact that their internally grounded *dispositions* or *capacities* to engage in the relevant life-functions are perfectly intact and undamaged. These organisms can be revived relatively easily, without being in any way repaired.[8]

---

[7] Feldman 1992a, p. 170.

[8] This case for the conclusion that O is not dead differs somewhat from Feldman's case. Speaking of organisms in suspended animation, he says, ". . . since they can return to life again later, it seems that they have not died." (1998: 819). Here, Feldman seems to suggest that the mere fact that something is such that it is *possible* for it to live later entails that the thing is not currently dead. I, on the other hand, am basing my conclusion that O is not dead on the fact that O's internally-grounded dispositions to engage in life-functions are fully intact and that O could be revived *so easily*, and without first being repaired. As to the

The relevance of these remarks to PD3 is easy to see. Consider some organism, O, that is now in suspended animation. Since O is not alive now but was alive at an earlier time, PD3 counts O as being dead. In fact, however, O is not dead; it is merely in suspended animation. This shows that PD3 is unacceptable. We need to define 'dead' in such a way that organisms in suspended animation do not count as dead. PD3 fails on this score.

2.3.2 The Problem of Deathless Fission

The second problem for PD3 arises from certain sorts of biological fission. According to Jay Rosenberg, when an amoeba undergoes a typical case of fission, it ceases to be present (and *a fortiori* ceases to be alive[9]), but it does not then *die*. Rather, it ceases to be present and ceases to be alive *deathlessly*. Here is Feldman discussing Rosenberg:

> Rosenberg asserts that death is not the only route out of life. To illustrate his point, he describes the case of an amoeba, Alvin [1983: 21-22]. He tells us that Alvin was a fat and healthy amoeba. According to the story, Alvin was so fat and healthy that at precisely midnight on Tuesday night/Wednesday morning, Alvin underwent fission and became two amoebas. According to Rosenberg, Alvin no longer existed [alternatively: no longer was present – CG] on Wednesday. Apparently, Alvin was "replaced" by his two descendants, Amos and Ambrose. Rosenberg claims that Alvin's example shows that "there are other ways for a life to come to an end besides death." [1983: 21] So while Alvin is no longer among the living on Wednesday, it is " . . . clear that he did not die." [1983: 22] My own intuitive sense of the situation is identical to Rosenberg's. I would not say that Alvin died. (1992: 66)

I share the Rosenberg-Feldman intuition and am inclined to agree that Alvin did not die when, at midnight Wednesday morning, he underwent fission and ceased to exist (or at least ceased to be present). But if this is correct, then it would be wrong to say that Alvin is *dead* on Wednesday afternoon. If Alvin made the transition from life to permanent non-existence (or non-presence)

---

suggestion that it is impossible for dead things to return to life, I find this highly implausible. Indeed, it seems to me that the Biblical story of Lazarus, in which Jesus brings a dead man back to life, is quite clearly possible in the metaphysical or broadly logical sense. (Tim Bayne uses this example to make the same point, MS.) Elsewhere (e.g., 1992b: 113), Feldman is very explicit about his wish to remain neutral as to the possibility of dead things returning to life.

[9] As Feldman has noted, the concept 'x is alive at time t' obeys the following principle: "necessarily, if a thing is alive at a time, then it exists [alternatively: is present – CG] at that time." Feldman claims, correctly in my opinion, that 'alive' is different in this respect from 'dead,' in light of the truth of the following principles: "possibly, a thing is dead at a time even though it does not exist [alternatively: is not present – CG] at that time," and "possibly, a thing is dead at a time even though it does exist [alternatively: is present –CG] at that time." (1992b: 114).

without ever dying in the process, then surely he doesn't get to count as *dead* at any time after the final moments of his career. This follows from a principle that I take to be self-evident: a necessary condition for being dead at a given time is having died at an earlier time. (I am not claiming that this condition is *sufficient*, however.)

To see why the case of amoebic fission poses a problem for PD3, consider some time t on Wednesday, after Alvin has divided. By time t, Alvin has permanently ceased to exist (or to be present) and *a fortiori* permanently ceased to be alive. But, as I just argued, Alvin is not dead at t. According to PD3, however, Alvin *is* dead then, since he is not alive then but was alive at some earlier time. This shows once again that PD3 is unsatisfactory. We need to define 'dead' in such a way that an amoeba does not count as being dead after it has divided into two new amoebas.

3 A Unified Solution to the Problems of Suspended Animation and Deathless Fission: Rosenberg's Definition

In light of the foregoing pair of problems, it may begin to appear that there is no satisfactory analysis of 'dead' that appeals just to 'alive' and standard logical and temporal concepts. Let us, therefore, consider a different approach. In particular, suppose that we have an adequate grasp of the concept 'x dies at time t.' Intuitively, to say that a thing *dies* at a time is just to say that it *begins* to be dead at that time.[10] Suppose, however, that our understanding of 'dies' does not depend upon any conceptually prior understanding of 'x is dead at t.' We could then appeal to 'dies' in an analysis of 'dead' in the following manner[11]:

D1 x is dead at t =df. x died at some time earlier than t.[12]

Indeed, D1 appears to capture Rosenberg's proposal. He writes that

[10] I intend this not as an official definition but merely as an informal reminder of the difference in meaning between 'x is dead at t' and 'x dies at t'.
[11] Of course, we would, on pain of circularity, be forced to abandon any attempt to define 'x dies at t' in terms of 'x is dead at t.' We could not, e.g., define 'x dies at t' as 'x begins to be dead at t.'
[12] Feldman (1992b: 108).

> . . . 'Aunt Ethel is dead' . . . seems to say just what 'Aunt Ethel has died' says . . . To say that a person is dead, then, seems . . . to report on a past event rather than a present condition . . . 'Being dead,' as we customarily speak, picks out only the "nominal condition" of having died. (1998: 42-43).

Although Rosenberg does not explicitly emphasize this point, his analysis of 'dead' has an important virtue: it furnishes us with a unified solution to the problems of suspended animation and deathless fission, the problems which forced us to abandon PD3.

We can begin with the problem of suspended animation. Return to the case of our organism O that is currently in suspended animation and hence neither alive nor dead. What does D1 have to say about O? According to D1, O is now dead just in case O died at some earlier time. Did O die at some earlier time? Presumably not. The only remotely plausible candidate for being a time at which O died is the time (call it *t1*) at which O entered suspended animation. Intuitively, however, it seems clear that O did *not* in fact die at t1: although O's life-functions then came to a halt and O ceased to be alive, O was not then damaged in any way, and it retained all of its internally grounded dispositions to engage in life-functions. Thus it appears that any acceptable account of 'x dies at t' will yield the result that our organism O did not die at t1.[13]

Assuming, then, that there is no time prior to the present at which O died, D1 entails that O is not currently dead, which is the result we want. In other words, D1 solves the problem of suspended animation. Since those things that are currently in suspended animation have not previously *died*, D1 entails, as it should, that they are not now *dead*.

Now that we've seen how Rosenberg's analysis of 'dead' solves the problem of suspended animation, let us apply this analysis to the problem of amoebic fission. According to PD3 (which says that a thing is dead at a time t iff it is not alive at t but is alive at an earlier time),

[13] It follows from this that we must not define 'dies' as follows: x dies at t =df. x ceases to be alive at t. Feldman discusses the problem of suspended animation in the context of his attempt to define 'dies.' After noting that the problem forces us to abandon the above definition of 'dies,' he argues, convincingly in my opinion, that there are other suspended animation cases that force us reject both of the following definitions (1) x dies at t =df. x ceases permanently to be alive at t, and (2) x dies at t =df. x ceases permanently and irreversibly to be alive at t. He concludes his discussion of the problem of suspended animation by suggesting, somewhat tentatively, that following definition may avoid that problem: "(D5) x dies at t =df. (i) ceases to be alive at or before t, and (ii) at t, internal changes occur in x that make it physically impossible for x ever to live again" (1992: 65).

Alvin the amoeba is currently dead, since he was alive until midnight last night, when he divided into two new amoebas, thereby permanently ceasing to be present and permanently ceasing to be alive. Rosenberg's analysis, however, does not force this implausible commitment upon us. According to his analysis, Alvin is now dead if and only if he died at some earlier time.

Did he? Well, there seems to be only one remotely plausible candidate for being a time at which Alvin died – viz., midnight last night, when he divided. As Rosenberg rightly noted, however, it seems false to say that Alvin died at that time. It appears, then, that any satisfactory account of 'dies' must yield the result that, in typical cases such as Alvin's, an amoeba does not die when it divides – although it does then permanently cease to be alive and permanently cease to be present.

So let us proceed on the assumption that there is no time (past, present, or future) at which Alvin dies. In that case, Rosenberg's analysis of 'dead' gives us the result we want – namely, that Alvin is not currently dead (even though he is no longer alive). If to be dead at a time is just to have died at some earlier time, then Alvin, never having died, is not now dead.

4 Feldman's Objection to Rosenberg's Definition

As we have just seen, Rosenberg's analysis solves the problems of suspended animation and amoebic fission. Shall we declare his analysis a success? That would be premature. For consider the following passage from Feldman:

> It may at first appear that we can define 'dead' straightforwardly by appeal to 'dies' as follows:
>
> D1 x is dead at t =df. x died at some time earlier than t.
>
> But suppose it is possible for something to live again after death; suppose that genuine revitalization is possible. In such a case, D1 would force us to say that the revived entity is at once living and dead. It would be living because it has been revitalized; it would be dead according to D1 because it died at some earlier time. This seems wrong to me. It seems to me that once it has been revived, such an organism would no longer be dead. It would then simply be alive again. (1992b: 108)

I accept Feldman's case against D1. Let us say that a thing is *reanimated* just in case it dies and sometime later is alive again, not dead. In my view, Feldman is right to insist that any satisfactory definition of 'dead' must leave open the possibility of reanimation.

There are two different reasons one might have for accepting this constraint on definitions of 'dead.' First, one might positively *endorse* the possibility of reanimation (as I do), either because one simply finds it plausible without argument, or because one thinks it can be argued for in a convincing fashion. Second, even if one does not actually endorse the given possibility, one might wish to avoid taking a stance on it one way or the other. Thus, one might be inclined to say, "Perhaps reanimation is possible, perhaps not. Whatever the correct answer may turn out to be, this answer should not follow directly from our definition of 'dead.' Any such definition should *leave open* the question of whether reanimation is possible." (This seems to be Feldman's reason for rejecting D1.)

5 Feldman's Definition

Having rejected D1, Feldman goes on to offer an alternative definition of 'dead' in terms of 'dies':

> I prefer to define 'dead' in this way:
>
> D2 x is dead at t =df. x died at some time earlier than t, and x has not been alive since then.
>
> So the concept of being dead can be defined by appeal to the concepts of dying ['x dies at time t'] and being alive ['x is alive at time t'], together with the concept expressed by 'earlier than.' (1992b: 108)

I agree that D2 avoids the problem that led us to reject its predecessor: unlike D1, D2 allows for the possibility of reanimation. To see this, suppose that John died in 1990, returned to life in the year 2000, and continues to be alive at the current time. In that case, D2 counts John as being dead at each time in the 1990s after his death, which is obviously correct. D2 does not, however, count him as being dead now, since it is no longer true that John "has not been alive since" the moment when he died. Again, this is the desired verdict.

Nevertheless, it seems to me that D2 is not quite right either. Before I explain why this is so, I would like to note that D2 is perhaps somewhat ambiguous.[14] Rather than pursue exegetical

[14] Some uncharitable interpretations of D2 are as follows:

issues in the main text, however, I shall simply assume that Feldman intends D2 to be understood in such a way that it is equivalent in meaning to:

D2* x is dead at t =df. there is some time t* such that (i) t* is earlier than t, (ii) x dies at t*, (iii) x is not alive at t or at any time between[15] t* and t.

Although several competing interpretations of D2 are available (see the foregoing note), it is fairly clear that D2* is both (i) a more plausible analysis of 'dead' than any of its competitors, and (ii) more likely than these competitors to capture to what Feldman had in mind.

6 An Objection to Feldman's Definition: Things In Suspended Animation That Were Dead More Recently Than They Were Alive

To see why D2 will ultimately have to be rejected, consider the following example. Suppose that, while I am on a tour of a biology laboratory, my guide shows me two frozen specimens, A and B. Specimen A is an organism in suspended animation. As I noted earlier, I follow philosophers such as Wreen and Feldman in holding that such organisms are neither alive nor dead.

Now suppose that my guide turns her attention to specimen B and supplies me with two pieces of information about it. First, B was once a living thing. Second, B is currently an intrinsic, molecule-for-molecule duplicate of A. Thus the two frozen specimens are in precisely the same overall intrinsic condition. I think it follows from this that specimen B, like specimen A, is

---

(F1) x is dead at t =df. there is some time t* such that: (i) t* is earlier than t, (ii) x died at t*, and (iii) x has not been alive since *t*.
(F2) x is dead at t =df. there is some time t* such that: (i) t* is earlier than t, (ii) x died at t*, and (iii) x has not been alive since *t**.
(F1*) x is dead *now*, at the *present* time =df. there is (or was) some time t* such that: (i) t* is in the past, (ii) x died at t*, and (iii) x has not been alive since the present time.
(F2*) x is dead *now*, at the *present* time =df. there is some time t* such that: (i) t* is in the past, (ii) x died at t*, and (iii) x has not been alive since t*.
Rather than subject the reader to the tedium of explaining why D2* is more plausible as definition of 'dead' than any of the above, I shall "leave this as an exercise."

[15] Temporal betweenness can be defined in terms of 'earlier than': time t2 is between time t1 and t3 =df. (i) t1 is earlier than t2 and (ii) t2 is earlier than t3.

currently in suspended animation.[16] Or, somewhat more cautiously, I think this conclusion follows *provided that A and B have been duplicates of each other for long enough, and will continue to be duplicates for long enough*.[17] Let us, then, adjust the example accordingly, so that A and B have been, and will continue to be, in their present state for as many years as you like. This gives us the conclusion that

(1) specimen B, like A, is in suspended animation.

And (1), together with the assumption that being in suspended animation is incompatible both with being alive and with being dead, entails that specimen B is neither alive nor dead at the present moment.

Now, to complete the example, suppose that my guide tells me that, whereas specimen A was alive just prior to entering suspended animation, *specimen B was dead just prior to entering this state!* In particular, B came into existence as a living thing, lived a normal life, and died. It continued to exist as typical dead thing for some time thereafter. Only then did its history take an unusual turn: a team of advanced biological engineers froze it, thereby halting its decomposition. While keeping it frozen at a very low temperature, the engineers began to alter its internal structure, molecule by molecule, until they had made it into an intrinsic duplicate of specimen A, in which condition it has been ever since. Thus B came into existence as a living thing, died, remained dead for some time, and then went directly from being dead to being in suspended animation, hence being neither alive nor dead. See the diagrams below.

A's timeline

[16] One principle that would license this inference is: necessarily, for any x and y and time t, if (i) x is in suspended animation at t, (ii) y is alive at some earlier time, and (iii) x and y are intrinsic duplicates at t, then y is in suspended animation at t. One potential problem with this principle is that it makes *x and y's being duplicates for a mere instant* sufficient for y's being in suspended animation, provided that y was once alive and that x is in suspended animation. See the next note for a principle that avoids this problem while still serving my purposes.

[17] Here I rely on the following: necessarily, for any x and y and time t, if (i) x is in suspended animation at t, (ii) y was alive at some earlier time, and (iii) x and y are intrinsic duplicates at t, then there is some finite number n such that, if x and y are intrinsic duplicates not only at t but also at each time in the preceding and following n-minute intervals, then y is in suspended animation at t.

A is alive.	A is alive.	A is in suspended animation.			A is in suspended animation (hence neither alive nor dead).

B's timeline

B is alive.	B dies.	B is dead.	B is frozen.	B is repaired.	B is a duplicate of A (hence neither alive nor dead).

I grant that this situation is surprising and bizarre, and perhaps even contrary to the laws of nature (though that's highly doubtful). I do not, however, think that this situation is impossible in the metaphysical or broadly logical sense. Hence I think that we should reject any definition of 'dead' that entails its impossibility. After all, if we want to avoid defining 'dead' in a way that rules out the possibility of things that go directly from being dead to being alive (and not dead), as both Feldman and I do, then surely we should also want to avoid defining this concept in a way that rules out the possibility of things, such as B, that go directly from being dead to being neither alive nor dead.

But D2 does rule out this latter possibility. Roughly put, it says that B is dead now, at time t, just in case B has died more recently than it has been alive. More precisely, D2 says that B is dead at t just in case there is a time t* such that: (i) B dies at t*, (ii) t* is earlier than t, and (iii) B is not alive at t or at any time between t* and t. Is there a time t* that meets these three conditions? Yes: the moment at which B died (call it "t-") is such a time. It is true that (i) t- is a time at which B dies. It is true that (ii) t- is earlier than t. And, since B is either dead or in suspended animation at each time between t- and t, including t, it is true that (iii) there is no time between t- and t at which B is alive. So D2 counts B as being dead now, at time t.

This is unacceptable. As we have seen, there is a persuasive argument to the effect that B is neither alive nor dead now. I conclude, therefore, that D2 must be rejected. We need to define 'dead' in such a way as to allow *not only* for the possibility of things that have gone directly from being dead to being alive, *but also* for the possibility of things that have gone directly from being dead to being neither alive nor dead.

If I were willing to appeal to the concept ‘x is in suspended animation at t’ as an undefined primitive at this point, I might be tempted to offer the following analysis of ‘dead’:

D3 x is dead at t =df. there is some time t* such that: (i) t* is earlier than t, (ii) x dies at t*, (iii) x is not alive at t or at any time between t* and t, and (iv) x is not in suspended animation at t.

This analysis would solve the problem about specimen B. Since B *is* in suspended animation at the present time, clause (iv) goes unsatisfied: analysis D3 counts B as not being dead, which is the result we want.

I think, however, that it would be a mistake to appeal to ‘x is in suspended animation at t’ as a primitive at this stage in our inquiry. If it were a clear, familiar, standard notion that many philosophers would need to appeal to *anyway*, regardless of any desire to define ‘dead’, then I would have no qualms about appealing to this notion. For in that case, such an appeal would not constitute an addition to our overall ‘ideology’, our total stock of primitive notions. As things actually stand, however, the notion ‘x is in suspended animation’ is *not* clear or standard or familiar, and it is not useful enough to have made a place for itself in the ideologies of most philosophers. In the name of clarity and ideological parsimony, therefore, I reject D3.

7 A New Definition of ‘Dead’: Being Dead As Requiring Possession Of A Toxic Property

I shall now begin the process of setting out a new analysis of ‘x is dead at t’, an analysis that is meant to avoid all of the problems we have discussed so far. The core idea underlying my proposal is this. In order for a thing to be dead, it is not enough that the thing have died and not come back to life. All of this is compatible with its being in suspended animation, and hence with the thing’s being neither alive nor dead. So, in addition to the requirements above, we must also demand that the thing have some property that is incompatible, not just with being alive, but also with being in suspended animation.

To see the force of this requirement, consider the property having a temperature of one degree centigrade above absolute zero in a world governed by the same laws of nature as @, (where "@" is a name for the actual world). Call this property *F*. F is incompatible with being alive: nothing can live at such a low temperature – not, anyway, in worlds governed by the same laws of nature as our own. But the property F would seem to be perfectly compatible with being in suspended animation. In order to be *dead*, a thing must have some property that is, so to speak, "more hostile to life" than the property F; to be dead, a thing must have some property that guarantees not only that its instances are *non-living*, but also that they are *not in suspended animation*. I shall call such properties *toxic*.

7.1 A First Attempt To Define Toxicity

In order to convert these rough suggestions into a fully satisfactory analysis of 'dead,' I will need to provide a definition of the technical term 'property F is toxic,' and I will need to do this without appealing to 'x is in suspended animation at t'. Here is a first attempt:

T1 F is toxic1 =df. (i) F is a property, and (ii) for any x and any time t, if x is alive at each instant in some interval that immediately precedes[18] t and x has F at t, then x dies at t.

In addition to appealing to 'dies' and certain standard temporal and logical notions, T1 also appeals to the notions 'F is a *property*' and 'x *has* (i.e., *instantiates*) property F at time t'. Like the temporal and logical notions to which it appeals, these latter two notions have already made a place for themselves in standard philosophical ideology: many philosophers have found it necessary to appeal to these notions in dealing with a wide variety of philosophical issues having nothing to do with the analysis of 'dead'. For this reason, we can appeal to them without incurring much cost.

[18] The idea here is that the interval is open toward the future and that t is its future-oriented end-point. In that case the interval does not include t, but there is no "gap" between the interval and t.

Let us pause for a moment to consider the notion of toxicity1 thus defined. To say that a property F is toxic1 is, roughly, to say that nothing ever *has* or ever *will* go directly from being alive to having F without dying immediately.[19] The goal is to capture the idea that a toxic property is one whose acquisition by a living thing immediately brings *death* (rather than mere *non-life*) to that thing. Consider, for example, the property being <u>being frozen</u>. Although no frozen things are alive, <u>being frozen</u> is not a toxic1 property, since many things have gone directly from being alive to being frozen without dying. <u>Being a molecule-for-molecule duplicate of the (long-dead) Lenin as he is at some moment in 2005</u>, however, is a toxic1 property, since nothing has ever gone directly from being alive to having that property without dying (and nothing ever will).

Now, to see how the notion of toxicity1 might be used to define 'dead', consider the following definition:

D4 x is dead at t =df. x dies at some time earlier than t, and x has some toxic1 property at t.

According to D4, the difference between (i) genuinely dead things and (ii) things that have died and then either come back to life or gone into a state of suspended animation is that the things in group (i) have some toxic1 property, whereas the things in group (ii) don't.

To see what is wrong with this suggestion, consider some particular bacterium, Bob, that is currently in suspended animation. (Bob is being stored at a very low temperature.) Since being put into this state, Bob's intrinsic properties have changed very little. They have changed slightly, however. Suppose that Bob's cell membrane suffered an extremely tiny indentation a few days after Bob was frozen; just two or three of Bob's constituent molecules were somewhat displaced.

[19] When I say "x goes from being alive to having F" I do not mean this to entail that the x ceases to be alive when it acquires F. It is meant to entail only that x is alive in some interval immediately prior to the first moment at which x has F; this leaves open the possibility that x continues to be alive even after it begins to have F.

If this happened, then in all likelihood Bob's current shape[20] is distinct from (though very similar to) any shape that Bob has had previously. Indeed, Bob's current precise shape is probably distinct from any shape that any living thing has ever had or ever will have.

But if so, then Bob's current shape is toxic1: nothing has ever gone directly from being alive to having the precise shape that Bob currently has. Certainly Bob himself did not do this; he did not go *directly* from being alive to having his current shape. Rather, some time passed between Bob's being frozen and his acquiring his current shape. Nor has anything else gone directly from being alive to having that shape. After all, no other organism has ever or will ever possess precisely this determinate shape. Since nothing ever has or will go directly from being alive to having the relevant shape, it's also true that nothing has ever done this *without dying immediately*. This means that Bob's shape qualifies as toxic1.

But now let us add to the story of Bob by supposing that, sometime prior to being put into suspended animation, Bob died. We need not concern ourselves with what happened between the time of his death and the time at which he enterer suspended animation. Perhaps he was brought back to life and lived again for a while, or perhaps he went directly from being dead to being in suspended animation. In any event, we know that Bob died, and we know that he currently has a toxic1 property. Thus we know that D4 counts him as being dead, and wrongly so. (Bob is currently in suspended animation and hence is not dead.) D4 must therefore be rejected.

The most obvious problem with D4 is that makes use of the wrong notion of toxicity. We wanted a notion of toxicity according to which all toxic properties are incompatible with being in suspended animation. As we have just seen, however, having a toxic**1** property is perfectly compatible with being in suspended animation. All that has to be true of a property F for it to count as toxic1 is that it be such that, *as a matter of actual fact*, nothing ever has or will go directly from being alive to having F without immediately dying. But there are many properties, such as certain precise shapes, that, *by sheer accident*, never have been and never will be

[20] By "shape" I mean "maximally precise, maximally determinate shape."

possessed by something that had been alive just prior to possessing the property. Many such properties perfectly well *could* be possessed by things in suspended animation and even by things that are alive; it's just that these properties *happen not to be* so possessed. Surely it is wrong to suppose that such properties are in the relevant sense "hostile to life."

7.2 Defining Toxicity via Necessity

Let us, therefore, specify a new notion of toxicity. Whether or not a given property is toxic1 depends entirely on how things stand in the actual world. Our new notion, toxicity2, will be different. Whether or not a given property is toxic2 will depend, not just on how things stand in the actual world, but also on how things stand in all other possible worlds. We will achieve this result by appealing to the notion of *necessity*:

T2 F is toxic2 =df. (i) F is a property, and (ii) *necessarily*, for any x and time t, if x is alive at each instant in some interval that immediately precedes t and x has F at t, then x dies at t.

The notion of necessity appealed to here, sometimes called metaphysical or broadly logical necessity, is invoked by a great many philosophers for an extremely wide variety of philosophical purposes. Our appeal to this notion in T2, therefore, should not be regarded as a significant addition to our overall ideology.

To say that a property F is toxic2 is, roughly, to say that it is *impossible* for a thing to go directly from being alive to having F without immediately dying. In other words, it is to say that the property is not just *actually* but *necessarily* toxic1. The current shape of our frozen bacterium, Bob, is toxic1 but only contingently so. Therefore that shape does not qualify as toxic2. (Thus toxicity2 entails toxicity1, but not *vice versa*.) Although nothing ever has or will go directly from being alive to having the shape in question, this is not a necessary truth: it is entirely *possible* for a thing do this.

On the other hand, it would seem that there are some properties that *absolutely could not* be acquired by a living thing unless the thing were to die at the moment of acquisition. Consider, for example, the property of being a molecule for molecule duplicate of V. I. Lenin as he was at midnight, 1 January 2000. Call this property being Lenin-like. It seems clear to me that it is impossible for a living thing to acquire this property without dying at the moment of acquisition.[21]

With the notion of toxicity2 in hand, we are in a position to consider our next analysis of 'dead':

D5 x is dead at t =df. x dies at some time earlier than t, and x has some toxic2 property at t.

Like D4 before it, D5 attempts to distinguish between (i) genuinely dead things and (ii) things that have died and then either come back to life or gone into a state of suspended animation. According to D5, the difference between these two groups is that the things in group (i) have some toxic2 property, whereas the things in group (ii) don't.

Unfortunately, D5 also fails. To see why, return to the case of Bob, our bacterium in suspended animation. Let us suppose that Bob has been frozen for the past three months. Thus it appears that Bob currently has the property of having been non-living for at least two months. But this property is toxic2! (It is impossible for anything to go directly from being alive to having

[21] Suppose that a living thing were to have this property for an extremely brief period of time, perhaps just for a single instant, then return to its previous condition. Some philosophers might claim that in this case, the living thing acquires the property without dying. I am inclined to disagree and to say that (if the situation is possible) the thing dies. But suppose for the sake of argument that my critic's claim is correct. Then being Lenin-like is not toxic2. Nevertheless, we can still produce a clear example of toxic2 property. Consider being Lenin-like for a long time to come, where it's a necessary truth that, for any x and time t, x has the relevant property at t if and only if x is Lenin-like at t *and at each instant in the five hours immediately following t.* It's clearly a necessary truth that if one is alive now but one is about to suddenly become an intrinsic duplicate of the long-dead Lenin *and one is to remain in that condition for five hours*, then one is about to die. (Being Lenin-like for a long time to come is, in Roderick Chisholm's terminology, 'rooted outside the time at which it is had.' Chisholm's attempt to define this notion (of which we seem to have an antecedent grasp) is as follows: "Necessarily, for any x and for any period of time t, x has the property G throughout t only if x exists at some time before or after t" (1976: 100). The fact that this definition actually fails to capture its target notion has spurred much of the literature attempting to define the related notion of intrinsicness. See Weatherson (?) for a survey.)

been non-living for at least two months, and *a fortiori* it is impossible for anything to do that *without dying immediately*.) Since Bob has died previously and currently has a toxic2 property, D5 counts him as being dead at the present time. This, of course, is unacceptable: Bob is currently in suspended animation, not dead. So D5 must be rejected.

Once again, we seem to have appealed to the wrong notion of toxicity. A property F can count as toxic2 merely by being such that it's impossible for anything to go directly from being alive to having F. But, as we've just seen, this won't guarantee that F is incompatible with being in suspended animation: having been non-living for at least two months is perfectly compatible with being in suspended animation, but it's toxic2 in the sense that nothing can go *directly* from being alive to having been non-living for at least two months without dying immediately. (Nothing can make such a direct transition *at all*, much less could anything do this without dying immediately.)

7.3 Defining Toxicity via Necessity and Temporal Locality

Our goal, again, is to define a notion of toxicity according to which toxicity guarantees incompatibility with being in suspended animation (as well as with being alive). To do this, it will be helpful to reflect on the problem that doomed our previous notion, toxicity2.

The problem with this notion, it seems to me, was that it was defined in such a way that it could apply to certain *temporally non-local* properties such as having been non-living for at least two months. Intuitively, to say that a property F is temporally non-local is to say that whether or not a thing has F at a time t depends *constitutively* (not merely *causally*) on how the world is before or after t. Clearly, whether or not a thing has the property having been non-living for a least two months at a time t is *not* purely a matter of how the world is at t; it's also a matter of how the world was throughout the two months leading up to t. By contrast, the property being three feet away from a piece of iron and the property having a mass of 3 kilograms both seem to

be *temporally local*:[22] for each of those properties, whether or not a thing has that property at a time t constitutively depends only on how the world is at t; what goes on at other times is relevant at most causally, never constitutively.

If we keep in mind the distinction between temporally local and temporally non-local properties, it becomes fairly easy to see why a property can qualify as toxic2 without being relevantly "hostile to life." Recall that a property F will count as toxic2 just in case it's impossible for a thing to go directly from being alive to having F at a time t without dying at t. But if F is a temporally non-local property, it may have *history-involving possession conditions* that by themselves make it impossible for anything to go directly from being alive to having F (and hence impossible for anything to do this without dying).

This can be so even when F, far from being *hostile* to life, is positively "life affirming". Consider the property being alive for the first time, and suppose that, necessarily, a thing has this property at an instant of time t if and only if the thing is alive at t but is not alive at any prior time. No one would accuse this property of being hostile to life! And yet it counts as toxic2: it's impossible for a thing to go directly from being alive to having that property (and hence impossible for a thing to do that without dying immediately). What accounts for the relevant impossibility is not hostility to life but rather the peculiar history-involving possession conditions of this temporally non-local property.

Accordingly, it appears that we can isolate a more relevant notion of toxicity by adding a *temporal locality* requirement to our previous definition. The resulting definition is this:

T3 F is toxic3 =df. (i) F is a temporally local property, and (ii) *necessarily*, for any x and time t, if x is alive at each instant in some interval that immediately precedes t and x has F at t, then x dies at t.

[22] Although both appear to be temporally local, only the mass property is intrinsic. To say that a property F is intrinsic is, very roughly and intuitively, to say (i) that F is temporally local and (ii) that whether or not a thing has F constitutively depends only on what that thing is like in itself and is at most causally dependent on how that thing is related to things outside itself. Clearly, being three feet away from a piece of iron is not intrinsic.

Admittedly, there is no widespread consensus about how to define 'temporally local,' and I will not attempt a definition here. Like the modal notion of *metaphysical necessity*, however, the notion of *temporal locality* has been invoked by many philosophers in a wide variety of contexts that are entirely independent of the attempt to define 'dead'. In short, it's a notion that many of us *already* need to depend upon. Thus I think that our appeal to it in the present context is justified. Of course, any obscurity in this notion will give rise to obscurity in the notion of toxicity3, but this is a risk that I am willing to take.

Now, what does T3 say? Roughly put, it says that a property F is toxic3 just in case F is temporally local and it's impossible for anything to go directly from being alive to having F without dying immediately. Since a property F can count as toxic3 only if F is temporally local, the impossibility of a thing's going directly from being alive to having F without dying immediately cannot be due to history-involving possession conditions (as was the case for some toxic2 properties). Therefore, it may begin to seem that the relevant impossibility could only be explained by F's *genuinely* being hostile to life in the sense we have been trying to capture.

Let us, then, consider a definition of 'dead' that appeals to the notion of toxicity3 we've just defined:

D6 x is dead at t =df. (i) x dies at some time earlier than t, and (ii) x has some toxic3 property at t.

According to D6, to be dead is, roughly, to have died at some earlier time and to have some temporally local property F that cannot be acquired by a living thing unless the living thing dies at the moment of acquisition. As before, a primary goal here is to distinguish between (i) genuinely dead things and (ii) 'undead' things, i.e., things that have died and then either come back to life or gone into a state of suspended animation. D6 tells us that the difference between these groups of things is that the dead things all have some toxic3 property, whereas the undead things don't.

7.4 Allowing For Non-Fatal Toxic Flickers

In my opinion, D6 is *far* superior to its predecessors. However, there still seems to be room for improvement. For I suspect that some philosophers will be moved by the following speech:[23]

> Clearly, Lenin is dead at midnight, 1 January 2000. Call that instant of time *tJ*. It is unlikely, however, that he has any toxic3 properties at tJ. Consider, for example, a property mentioned earlier: being Lenin-like, where, necessarily, a thing has this property at a time t iff that thing is a molecule-for-molecule duplicate of V. I. Lenin as he is at tJ. This property seems to be not only *temporally local* but also *intrinsic* in the sense described above. Moreover, if Lenin has *any* toxic3 properties at tJ, it would seem that being Lenin-like is one of them. (If this property isn't toxic3, why would any of Lenin's *extrinsic*, temporally local properties be toxic3?)
>
> Upon close examination, however, being Lenin-like ends up failing the test for toxicity3. Suppose that, while sleeping peacefully through an otherwise ordinary night in mid-2005, George W. Bush abruptly acquires the property being Lenin-like for an extremely brief period of time (perhaps just for a single instant, perhaps for a fraction of a nanosecond). Then, just as abruptly, he loses this property and returns to the condition he would have been in had things proceeded normally all along. Do we want to say that Bush *dies* when he acquires the property in question? The most reasonable answer seems to be No: Bush does not die at this moment. (If he had remained Lenin-like for a longer period of time, it would have been correct to say that he died, but since he returns to returns to his normal condition so quickly, he ends up not dying at all that night.) Thus the acquisition of being Lenin-like does not necessarily bring immediate death to a living thing, and this property does not count as toxic3. Presumably, then, none of Lenin's properties at tJ are toxic3. Thus D6 wrongly counts Lenin as not being dead at tJ.

My imagined objector is confident in the metaphysical possibility of Bush's flickering in and out of a Lenin-like condition. I do not share this confidence. And when I waive this concern and ask myself whether Bush would die were he to "flicker" in this way, I'm somewhat inclined to say: Yes, he would die (and then quickly return to life). So I lean toward the judgment that being Lenin-like is toxic3.

Other things being equal, however, I would like my definition of 'dead' to be as ecumenical as possible; and in particular, I would like this definition to be acceptable even to those who agree with the speech above. Thus I suggest that we abandon our notion of toxicity3 in favor of a notion defined as follows:

T4 F is toxic4 =df. (i) F is a temporally local property, and (ii) there is some positive real number n such that, necessarily, for any x and time t, if x is alive at each instant in some interval immediately preceding t, and x has F at t and at each instant in the n-minute-long interval immediately following t, then x dies at t.

[23] I borrow the useful device of the 'speech' from Peter van Inwagen.

Very crudely put, a property F is toxic4 provided that F is temporally local and it's impossible for a thing to go directly from being alive to having F for a "long enough" stretch of time without dying when it begins to have F.

My imagined objector, who denied that being Lenin-like was toxic3, will presumably agree that this property is toxic4. Suppose that Bush can flicker in and out of being Lenin-like without dying. Even so, it's clearly a necessary truth that if Bush becomes Lenin-like *and stays that way for long enough*, he will count as having died when he first began to be Lenin-like. Thus being Lenin-like is toxic4.

7.5 Denying Toxicity to "Super-Cold" Suspended Animation States

Consider, on the other hand, some organism O that has just ceased to be alive and has entered suspended animation. Now let F be O's current overall intrinsic (and temporally local) condition – i.e., the conjunction of all of O's current intrinsic properties. Is F toxic4? I think not. In the first place, it's pretty clearly possible for a thing to go directly from being alive to having F without dying: this is exactly what happens to O. Secondly, it seems that *no matter how long O continues to have F*, it will still remain true that O did not die when it first acquired F. If I am right about all of this, then F turns out to be non-toxic4, as it should. (F is not hostile to life in the relevant sense.) So far, so good.

Suppose, however, that we alter the case slightly. Instead of considering an organism that has just ceased to be alive, and whose current condition is therefore obviously compatible with *having just been alive*, consider an organism O* in suspended animation that ceased to be alive *several hours ago* and whose current temperature is very close to absolute zero. Let F* be O*'s current overall intrinsic condition. Is F* toxic4?

Unfortunately, it probably is. In order for a thing to be alive, the thing's constituent particles must be organized in the appropriate ways and must be undergoing the right sorts of activities. Anything that has F*, however, will be made up of particles whose organization and

activities are entirely *inappropriate* for being alive. Nothing that has F* will bear a very close resemblance (at the cellular and molecular level) to any possible living thing. Thus, in order for a thing to go *directly* from being alive to having F*, the thing would need to be capable of persisting through a radical and abrupt *discontinuity* with respect to the organization and activities of its constituent particles.

Like many other philosophers, I have grave doubts about the possibility of a thing's persisting through such a discontinuity. In other words, I suspect that it's impossible for a thing to go directly from being alive to having F* and hence impossible for a thing to do this *without dying immediately*. If my suspicion is correct, then F* qualifies as toxic4. But of course this is not what we want: since F* is the overall intrinsic condition of a thing *in suspended animation*, F* is not in the relevant sense hostile to life.

Fortunately, an improved notion of toxicity is close at hand. Indeed, I think we are now ready to define our *final* notion of toxicity:

T5 F is toxic5 =df. (i) F is a temporally local property, and (ii) there is some positive real number n such that, necessarily, for any x and times t and t*, if

(a) x is alive at t,

(b) t is earlier than t*,

(c) x has F at t*, and

(d) x has F at each instant in the n-minute-long interval immediately following t*,

then

(e) x dies at t* or at some time between t and t*.

Our previous notion of toxicity applied to a temporally local property F just in case it was impossible for a thing to go from being alive to having F (or having F for a long enough time) without dying *at the very moment the thing acquired F*. On the other hand, a temporally local property F counts as *toxic5* just in case it's impossible for a thing to be alive at *one* time, t, and to

have F at a *later* time, t*, and for a long enough stretch of time following t*, without dying at t* *or at some time between t and t**. As I will now explain, being Lenin-like is toxic5 but the overall intrinsic condition F* of our super-cold organism O* is not toxic5. In both cases, this is the result we want.

Let me begin with the organism O* in condition F*. This organism, recall, is being stored in suspended animation at a temperature of just above absolute zero. The question we need to ask is this: is it possible for a thing to be (i) alive at one time, t, and (ii) in F* at a later time, t*, (iii) in F* for as long as one likes following t*, *without dying at t* or at any time between t and t**? The correct answer is very clearly: Yes, this is possible. Recall that O* itself made the transition from being alive to having F* without dying. (Once O* ceased to be alive, however, some time passed before O* became cold enough to have F*. Thus the transition was not direct.) Now that O* has F*, it seems that O* can continue to have F* for any arbitrarily long period of time without ever dying: as long O* continues to have F*, O* will be an organism in suspended animation that has ceased to live but has not yet died. So F* is not toxic5.

What about being Lenin-like? Could a thing be (i) alive at one time, t, (ii) Lenin-like at a later time, t*, and (iii) Lenin-like throughout a "long enough" interval immediately following t*, all without dying at t* or any time between t and t*? I think not. I find it intuitively clear that, as a matter of necessity, if a thing is alive at one time then Lenin-like throughout a long enough interval later on, the thing dies sometime in between (perhaps upon first becoming Lenin-like). This is the sense in which being Lenin-like is hostile to life. Indeed, I think that hostility to life *just is* toxicity5.

The natural thought at this point is that we can appeal to the notion of toxicity5 in the following definition of 'dead':

D7 x is dead at t =df. (i) x dies at some time earlier than t, and (ii) x has some toxic5 property at t.

Initially, this may seem right: to be dead is to have died and to have a property that's hostile to life in the sense specified by our definition of toxicity5. But I think we need to make two changes to D7 before we reach a definition we can be happy with.

7.6 Being Dead As Requiring Possession Of A Toxic Property or Nonpresence

The need for the first change can be seen by returning to the case of Plato. According to D7, Plato is dead at the current time (call it tC) just in case he (i) died at some earlier time and (ii) has some toxic5 property at tC. Clearly, the first clause is satisfied: Plato did die at some earlier time. But what about the second clause? Does Plato have any toxic5 properties at tC?

Some philosophers will endorse the following argument for the conclusion that Plato does not have any toxic5 properties at tC:

1. Plato is not present at tC. (Perhaps he ceased to be present at the moment of his death; perhaps he continued to be present for some time (centuries even) after his death as a corpse. In any event, he's not present now.)
2. Necessarily, for any x and time t, if x is not present at t, then x does not have any (toxic5 or non-toxic5) properties at t.
3. Therefore, Plato does not have any (toxic5 or non-toxic5) properties at tC.

To be sure, this argument is controversial. In particular, premise 2 is likely to be challenged by some philosophers. Nevertheless, I suspect that many others will accept the argument and will agree that Plato has no toxic5 properties at tC. But when this conclusion is combined with D7, it yields the absurd result that Plato is not currently dead. So anyone who accepts the above argument will need to reject D7.

In the interests of accommodating these philosophers, therefore, I suggest that we replace D7 with:

D8 x is dead at t =df. (i) x dies at some time earlier than t, and (ii) either x has some toxic5 property at t or x is not present at t.

According to D8, both Lenin and Plato are dead at the current time, tC, even though only the former is *present* at tC. Lenin qualifies as dead by virtue of having died at an earlier time and currently having some toxic5 property; Plato qualifies by virtue of having died and failing to be present altogether at tC. So much for the *first* of the two necessary changes mentioned above.

7.7 Being Dead As Requiring A *Death* As One's Most Recent Exit From Potential Life

I turn now to the second change. Suppose that an amoeba, Larry, dies at time t1 and is revitalized at the later time t2. Larry remains alive from t2 to t3, at which time it undergoes fission (in the manner characteristic of amoebas) and ceases to be present. All times thereafter, including t4, are entirely devoid of Larry.

t1	t2		t3	t4
Larry dies.	Larry is revitalized.	Larry is alive.	Larry divides and ceases to be present (permanently).	Is Larry dead now?

It is easy to see that D8 counts Larry as being dead at t4. According to D8, Larry is dead at t4 just in case (i) Larry dies at some time earlier than t4 and (ii) Larry either has some toxic5 property at t4 or fails to be present at t4. And there is no doubt that both of these clauses are satisfied. Since Larry dies at t1, clause (i) is satisfied; and since Larry is not present at t4, clause (ii) is satisfied. According to D8, then, Larry is dead at t4.

But if we return to our earlier discussion of amoebic fission, we can see that there is a forceful case to be made for the conclusion that Larry, rather than being dead at t4, is in fact *neither alive nor dead* at t4. Rosenberg and Feldman argued, correctly I think, that when an amoeba undergoes a typical case of fission, it ceases to be present (and *a fortiori* ceases to be alive), but it does not then *die*. Rather, it ceases to be present and ceases to be alive *deathlessly*.

Following Rosenberg and Feldman, then, shall suppose that Larry does not die at t3, when he undergoes fission. Although he does then cease to be alive and cease to be present, he does so deathlessly.

In light of this, surely it would be wrong to say that Larry is *dead* at t4 or at any time after the final moments of his existence. However, this is exactly what D8 says: as we saw earlier, D8 tells us that Larry is dead at t4. Thus I suggest that we abandon D8 in favor of:

D9 x is dead at t =df.

(i) Either x is not present at t or x has a toxic5 property at t, and

(ii) there is some time t* such that:

(a) t* is earlier than t,

(b) x dies at t*, and

(c) for each time tB between t and t*, either x is not present at tB or x has a toxic5 property at tB.

This definition yields the desired result that Larry is not dead at t4. According to D9, Larry is dead at t4 just in case: (1) Larry either has a toxic5 property at t4 or is not present at t4, and (2) there is some time t* with the following features: (A) t* is earlier than t4, (B) Larry dies at t*, and (C) Larry is either fails to be present or has a toxic5 property at each time between t* and t4. Is there a time with those features? No. There is only one time at which Larry dies – namely, t1. Although t1 is earlier than t4, it's not true that Larry either fails to be present or has a toxic5 property at each time between t1 and t4. After all, Larry is actually *alive* at each instant between t2 and t3, and all of these times are between t1 and t4. So D9 entails, as it should, that Larry is not dead at t4. Where D8 fails, D9 succeeds.

8 Conclusion

D9 is satisfactory as a definition of 'dead' only if we have an adequate grasp of each of the notions to which it appeals. Aside from purely logical notions, they are: (1) 'x *dies* at t,' (2) 'x is

alive at t,' (3) 'time t is *earlier* than time t*,' (4) 'T is an *n-minute-long interval* of time,' (5) 'x is *present at* t,' (6) 'F is a *temporally local property*,' (7) 'x *has* (instantiates) F at t,' and (8) 'it is *necessarily* the case that p.' If any of these notions turns out to be unintelligible, then D9 will have to be abandoned. And if any of the above notions turns out to be intelligible *only if defined in terms of 'dead,'* then D9 will be circular.

I am inclined to think that we do have an adequate understanding of notions (2) – (8) and that our grasp of these notions does not depend upon any prior grasp of 'dead'. I am uncertain, however, that our grasp of 'dies' is prior to our grasp of 'dead.' Perhaps the concept 'x dies at t' must be defined in terms of 'x is dead at t,' as follows:

D10 x dies at t =df. x begins to be dead at t (i.e., there is some interval of time T leading up to t such that x is *not* dead at any instant in T, and there is some interval of time T* immediately following t such that x *is* dead at each instant in T*).

As I mentioned at the beginning of the paper, however, I take no stand on the question of whether D9 is ultimately satisfactory. My claim is merely that D9 is the best available definition of 'dead' in terms of 'dies'.[24]

References

Bayne, Tim. Unpublished manuscript.

Chisholm, Roderick M. 1976. *Person and Object: A Metaphysical Study* (La Salle, IL: Open Court).

[24] Throughout this paper I have ignored the possibility of worlds in which things travel backward in time – either by tracing out closed timelike curves or in some other manner (on backward time travel see Lewis 1986: 67-80, Savitt 1994, Yourgrau 1999, and Sider 2001: 101-109). If there are any such possible worlds, I suspect that they would supply counterexamples to D9. Suppose, e.g., that Lenin's corpse = Lenin, and that the corpse is transported to the year 1776. Then it may seem correct to say that Lenin is dead at some time t in 1776, even though there is no *earlier* time at which Lenin died. However, I also suspect that D9 could be modified in a way that salvages much of its structure and basic insights. Perhaps this could be done by replacing (i) all talk of *instants and intervals of time* in D9 with talk of *instants and intervals of proper (or 'personal') time* and (ii) all talk of the earlier than relation in D9 with talk of the earlier than with respect to proper time relation. (See Lewis 1986: 67-80 for a discussion of personal time.) In our imagined time-travel case, although there is no instant t* of (external) time earlier than t such that Lenin died at t*, it's still true that Lenin died at some instant tD of his proper time such that tD is earlier with respect to proper time than any instant tE of his proper time such that he's dead at tE and present in 1776 at tE.

Elder, Crawford. 2003. *Real Natures and Familiar Objects* (Cambridge: MIT).
Feldman, Fred 2000. "The Termination Thesis," *Midwest Studies in Philosophy* XXIV: 98-115.
Feldman, Fred 1998. "Death," in *Routledge Encyclopedia of Philosophy*, ed. Edward Craig (London: Routledge), 1998: 817-823.
Feldman, Fred 1992a. *Confrontations with the Reaper: A Philosophical Study of the Nature and Value of Death* (Oxford: OUP).
Feldman, Fred 1992b. "The Enigma of Death" *Philosophia* 21: 163-181.
Heller, Mark. 1990. *The Ontology of Physical Objects: Four-dimensional Hunks of Matter* (Cambridge: Cambridge University Press)
Hershenov, David 2005. "Do Dead Bodies Pose a Problem for Biological Approaches to Personal Identity?" *Mind* ???
Hoffman, Joshua and Gary Rosenkrantz 1997. *Substance: Its Nature and Existence* (London: Routledge).
Lewis, David. 1983. *Philosophical Papers Volume I* (Oxford: Oxford University Press)
Mackie, David 1999. "Personal Identity and Dead People," *Philosophical Studies* 95: 219-242.
Olson, Eric 2004. "Animalism and the Corpse Problem," *Australasian Journal of Philosophy* ???
Quine, Willard van Orman 1960. *Word and Object* (Cambride: Harvard University Press)
Rosenberg, Jay F. 1998. *Thinking Clearly About Death*, 2nd Ed. (Indianapolis: Hackett).
Savitt, Steven 1994. 'The Replacement of Time' *Australasian Journal of Philosophy*
Sider, Theodore 2001. *Four Dimensionalism: An Ontology of Persistence and Time* (Oxford: OUP).
Thomson, Judith 1997. paper in *Reading Parfit*
van Inwagen, Peter 1990. *Material Beings* (Ithaca: Cornell).
Weatherson, Brian. SEP entry on 'Intrinsic.'
Wreen, Michael J. 1987. "The Definition of Death," *Public Affairs Quarterly*, 1, 4: 87-99.
Yourgau, Palle 2000. "Can the Dead Really Be Buried?" Midwest Studies in Philosophy XXIV: 46-68
Yourgrau, Palle 1999. *Gödel Meets Einstein: Time Travel in the Gödel Universe*, Open Court
Yourgau, Palle 1987 "The Dead," *The Journal of Philosophy* 86 ???.