

Mind and Body: Mind *is* Body

Descartes bequeathed us a clear statement of what was at least implicit for centuries if not millennia before him. Using the Medieval/Aristotelian framework of substance and accident, he identified two substances – *res extensa*, extended substance, and *res cogitans*, thinking substance – launching, thereby, the modern version of the Western intellectual mind-body problem.

I don't know that this problem will ever be resolved to everyone's satisfaction, though it's not inconceivable. There are lots of variations on several themes on offer from various retailers of ideas. Some variations are materialist (i.e. based on there only being physically detectable things and events in the universe); others are prepared to accept non-material things as existing either in our universe or in some spiritual or ideal or otherwise non-material complement to our universe. I am inclined to be a materialist myself, but I'm sorry to have to say I think the various forms of firmly committed materialist offerings all suffer from a credibility flaw that other offerings avoid.

The flaw is this: To date, to my way of thinking, there have been no materialist explanations of how it is with us humans that have come close to telling a compelling story about how atoms or subatomic particles or what have you, all crashing around in or emanating vibrations through the physical, extended space could possibly create feelings like we each experience. I don't mean we each experience all the same feelings, I mean that perhaps

the most personally significant experiences we each have are the feelings that course through us, creating the harmonic atmospheres, the textures, the subtle not-quite-smelt perfumes, the almost-tasted flavors, the washes of nearly-color in the midst of which we live and act and, usually most compellingly experience ourselves as being ourselves.

This is the realm of what the philosopher David Chalmers dubbed ‘the hard problem’.

Whatever the merits of Chalmers’s reasoning on that topic – I’m not his biggest fan but I don’t want to get into the back-and-forth of philosophical debate on this subject – I think he did us a significant service by putting a label to the hook we really can’t squirm off of.

Because this hook is very important. What the materialists are saying to the rest of humanity is, in effect, don’t worry about what’s most important to you. We’ll figure it out eventually. In the meantime, let me show you this really cool thing we *have* figured out.

Then we get the speaker’s favorite new neural gee-whiz finding. Oh, and those feelings you have, about yourself and about everything else, aren’t really all that important. They’re fictions, you know. We don’t have to understand the ego any more urgently than we need to understand Mickey Mouse.

On the other side, the not-so-materialists are basically saying, in these terms, we may not understand the ego, but we sure believe in its existence and its central importance.

If it's your life that's on the line – and in these conversations I think we all do well to conceive of them in such terms – which party would you cast your lot with?

Well, however you might answer that question, I certainly feel the tug of the not-so-materialist perspective on the importance, the value of my sense of myself.

But I think this problem is due to a failure of imagination on the part of the firmly committed materialists. I think there *is* a plausible story to be told about how it's possible for insensate matter to have the feelings I have, the feelings you have. I think, though, that before that story can get told in sufficient detail to be scientifically compelling there is an enormous range of scientific disciplines whose results will need to be marshalled, not to mention new ideas that will need to be developed, and moreover in marshalling those results and developing those ideas numerous specific beliefs we have today will end up having to be changed. That's but the nature of scientific progress and it's not my ambition today to anticipate just what that outcome will be.

It *is* my ambition, though, to tell a story that suggests such an undertaking has a decent chance of success, that it will *not* result in concluding that human consciousness is an illusion, irrelevant to actual, real human interests, nor in the superseding of human control of human affairs by some higher intelligence. In fact, while the fully developed scientific telling of such a story (which I have no dreams of doing here) will undoubtedly provide us with enormous ranges of new capabilities for mitigating forms of human suffering we don't

currently know how to reliably alleviate, it will basically leave us just where we are in managing our affairs.

That last point is very important and needs a bit of explanation. I believe that decisions – making which and conforming our actions to which comprise what I mean by ‘managing our affairs’ – come in basically two flavors. There are decisions about what to do that can be resolved by the direct, explicit application of already established scientific or technological capabilities. These decisions are really just objective questions that need answers and the decision-maker’s success in achieving her or his purposes depends upon adequate answers and a reasonable willingness to conform her or his actions to the dictates of those answers.

Those are basically pretty straightforward. They may not be easy – the answers may require pretty complex, abstruse, and extensive research, calculation and testing during which lots can go astray, and wrong answers may have dangerously harmful consequences. But they at least offer the prospect of a reasonable degree of confidence in the outcome, provided the intellectual work is done adequately and the execution conforms to the strictures the intellectual work calls for.

The other kind of decision involves situations in which the best science we can apply to the question at hand fails to tell us what to do. We are then left with exercising our own judgment. That’s hard enough when we’re confronted with such decisions as individuals,

it's only the more complex, uncertain and difficult when many people are engaged in the decision-making and stand to gain or lose in the outcome.

I think the development of a scientifically grounded story about what it is for us humans to be conscious may well add to the range of topics that can be addressed with the first kind of decisions, i.e. decisions that depend exclusively upon questions which can be answered by appeal to expertise. But what it won't do, I believe, is eliminate the second kind of decision, decisions that require judgment in addition to or even without appeal to expertise. In fact, I don't think it'll really reduce their frequency.

I say this for several reasons. First, my experience with technology (and I spent my career as first a programmer, then a software designer, then a manager and finally a corporate systems executive) is that for every topic that it moves from the second kind of decision to the first, it creates new opportunities for decisions of the second kind. I think this happens in part because no decision of the first kind is ever really isolated from decisions of the second kind. There are reasons for this, an important class being pertinent to the process of developing and applying expertise to a question. For example, groups of people may often be faced with questions about how much of which experts' time should be allocated to any one decision, given that at any point in time several decisions can be pending.

Second, new technologies that substitute machinery for human action (whether physical or cognitive) also commonly create the need for processes outside the decision makers'

ken. Newly deployed technology can require added infrastructure capabilities, ranging from things as straightforward as additional electrical capacity to much more complex and potentially ambiguous things like complex maintenance and other operational services.

But third and I believe most significantly, new technologies commonly create new opportunities for action, actions that might have been impossible or impractical previously. Now that you can do something you could never do before there are new decisions to consider. Those decisions can be of both kinds, but since they are about *new* opportunities for action it's likely, especially early on in the adoption of a new technology, for there to be no expertise worthy of the name that's relevant to such decisions.

The point to all this is simply that we should not hope that any technology we invent is likely to displace the need for decisions of the second kind. This is what I mean by saying that new science and technology will only leave us right where we are in managing our affairs.

And that point is important because while one might imagine that, in principle anyway, decisions of the first kind might be settled by purely mechanical means (i.e. if our technologies continue to develop in the right way), decisions of the second kind are *not* amenable to mechanical means. They involve conscious human beings taking express responsibility for judgment calls at most partially underpinned by objective fact.

There's one way one might imagine that a sufficiently sophisticated technology could change that situation: by changing all decisions of the second kind into decisions of the first kind. Personally, I am deeply skeptical that this is anything near as likely as purveyors of science-fiction scare stories try to get us to accept. The universe is not limited by our capabilities, however vast those may be.

I might be wrong about this, but I don't see any way to settle the question. Rather than settle it, I will simply set it aside for purposes of the story I want to tell and accept that even if the story works today and for some period of time into the future, the time may come when it no longer does. But I sort of think that's the best anyone can hope for with any story. Moreover, I think that once we consider what a resolutely materialist story about our humanity could look like, one that does not set aside our consciousness and feelings and other 'intentional' phenomena as mere illusions, we'll see that it's in the nature of being the conscious beings we are that we can't make something that is a better form of us. We're stuck with ourselves pretty much as we are.

I also think that such a resolutely materialist story, one that has no illusions about its ability to supplant our responsibility with a provably reliable, mechanically dependable method for deciding what to do, will give a good reason to believe we can't do that. Or rather, that we can only supplant our responsibility with mechanical means by abdicating our responsibility, not fulfilling it more effectively.

I said earlier that I thought that the lack of a meaningful, resolutely materialist story about ourselves is due to a failure of imagination. I should qualify that a bit – it's not *only* a failure of imagination. During the last 50 years or so the improvement of our understanding of how our bodies actually work, including especially the stunning advances in cognitive neuroscience in the latter part of that period, have provided material for our imaginations that could not have existed before that time in anything like the durable forms it is now taking. But now that we have as much as we do, I think it's time to really think about how it's possible that material things can be conscious, about how such a story could even be faintly plausible.

What follows, then, isn't a hypothesis about the neuroscience of human consciousness. It is, instead, what I've been calling it – a story. I have the temerity to put this story forward because I think that it contributes to dissipating the otherwise highly understandable anxiety about the mindless zombie prospects that pure materialism's stories about us seem to waive in front of us.

I hope, in telling it, to contribute to reducing the heat surrounding the quest for such understanding. What's not in the story is any sense in which there is an eternal life, nor, indeed, a life after death of any sort. Accepting the story, then, requires accepting one's life is, or at least may be the only life one gets.

What the story does, in my view, is give us reason to accept that a resolutely materialist understanding of human life doesn't require that such a life be meaningless. Instead, it calls for us to recognize the meaningfulness of the life we get, meaningful even if limited only to our time here, brief and vanishing as it is. The only immortality we have is that we will always have been what we are. Even if no one remembers a thing about us, even if there is no one around any longer to remember anything, it will still be true that we existed as we are. That will still be true through all eternity. It doesn't give comfort about suffering and death, but I'm not sure that comfort is all it's cracked up to be anyway and in any case it's still a pretty astounding fact. I exist, you exist, we are having the experiences we have and I think this is an absolutely astounding fact and one which no scientific understanding can ever eradicate or even jeopardize.

A scientific understanding that simply accepts that the existence we each have cannot be eradicated except by violence, disease or decay, and accepts that even though each of our existences *will* be eradicated by one of those at some time to come, nonetheless that existence is still immensely valuable and meaningful, such an understanding need not be at war with understandings that hope for some extension beyond those mortal bounds. It can just disagree about that but still attend to a fuller understanding of how it is with us now while we're still alive.

Such an understanding invites us, though it doesn't require us, to abandon dreams of some much grander and more glorious authority or authorities who know all and can effect

whatever they want to effect and in the fulfillment of whose purposes our lives get their meaning, and substitute for that the recognition that the eternal existence that is the indestructible fact of our past (and current) existence has the potential to create every bit as much meaning as our lives need.

In very simple outline, the story is this: selfhood is something we all learn to have from our parents and others among whom we grow up. It is grounded in our biological natures as beings whose existence is sustained by physical motions that often involve our whole bodies, motions that must be coordinated to the end of acquiring access to, consuming and transforming the resources necessary for our continued existence. But it's something we learn, just like we learn language – by which I mean something the learning of which is perhaps better thought of as a developmental process (or perhaps several such) that creates structures in our bodies (neurological but not only neurological) that will not change further as much as they have changed since birth (like our bodies reach a mature size and no longer get bigger in certain important structural respects). The results are things we can't unlearn or forget, just like we literate native English speakers can't unlearn or forget that the shapes inside the braces: {TREE} mean tree. I flesh that out more below, but first I add a bit of background in the hope of avoiding misunderstanding.

Many other beings, all sorts of animals in particular though perhaps even some plants, can be said to have similar capabilities to move themselves about. The differences in this regard between us and what we know of other beings' existence, differences which I'm

content to imagine may end up being truly small variations from capabilities some perhaps even many other sorts of beings have, isn't the awareness of our bodies' orientation in the world nor the ability to choose an action, initiate it and adjust its trajectory to the fulfillment of the purpose for which it was initiated (getting the forkful of food from the plate to our mouths as it were), the differences are in our awareness of ourselves and those around us as being aware, as being a unique instance of an existence each one of us shares.

We can presume, for example, that a dog can smell a rose every bit as specifically and clearly as we can, indeed probably with shockingly more accuracy and nuance than we can muster by way of our much more limited noses. What we shouldn't presume, though, is that the dog is capable of imagining how other dogs might react to that rose, what significance it might have for them, what, in other words, it might be for the rose to be a rose by that name or another.

It's that awareness of our experiences as the experiences others of us might also have that is the way I like to think of the core of consciousness, the presence of qualia as some might put it, the sense of being me. Every animal that moves itself about the world must embody the capability of coordinating the motions of its bodily parts so that the motion is effected and so that the effected motion contributes to acquiring, or at the very least doesn't squander too much of the resources it must have to survive. That degree of coordination may or may not be instantiated (in its nervous system, e.g.) in a process that we would recognize as being some consolidated awareness of itself and its purposes, that I suspect

would depend upon just how much the animal's survival did in fact depend upon it having such consolidated awareness. But it must work in a way that we could see, from the outside, as serving the purpose of sustaining its life. Otherwise it would be dead and not moving.

Every animal, that is, must respond successfully to things in the environment around it, things which in many cases we human observers can see in its environment because they're experientially in (or capable of being in) our environment too. ("Success" here means contributing to or at least not dangerously impairing the animal's potential for survival.) But not every animal must have such self-awareness as we experience in order to pull off that trick. I'm not claiming that no animal but a human animal *can* or *does* have self-awareness, I'm only claiming that it need not have self-awareness in our form to be able to successfully navigate the world. I'm not making an empirical claim here, remember, I'm telling a story and in this story, I'm simply not deciding the question of what other animals do. I'm really only invoking the fact that biological organisms have means of coordinating the motions of the parts of their bodies to the fulfillment of survival imperatives – a pretty obvious and humdrum point in its way – to say that our self-awareness exists in that context. We're another such animal and whether self-aware or not we have such abilities.

Now for an important bit of hand waiving: I believe that up to at least some level of complexity one can tell a resolutely materialist story about animals' existence. We needn't

invoke self-awareness and risk the mind/body problem to explain the existence of, say, earth worms or cockroaches or other creatures of comparable complexity. For purposes of my story I'm going to presume readers who go further are willing to let me make that assertion at least pro tem, as a basis for the story, deferring possibly needed revision to another occasion.

And a bit more energetic waive: I invite further accepting that quite a ways further 'up' the 'tree of life' from insects and invertebrates we can still say that we don't need to invoke a sense of self-awareness, just an ever more elaborate and sophisticated form of coordination. (Scare quotes around 'up' and 'tree of life' to indicate that I don't want to place any weight on the metaphors those words invoke other than a rough and not necessarily consistently contoured gradient of complexity and sophistication.) At some point, either with the emergence of the modern human species or somewhat earlier, self-awareness emerged. Up to that point, I'm saying (and this is the full-on hand waiving), we can presume there is a resolutely materialist explanation of how such beings can exist and, moreover, that we can presume such materialist explanations apply to our beings at least up to the point where questions of self-awareness become relevant to the explanation.

So the question my story is intended to answer, or rather to point at an illustrative example for the purpose of asserting that there *is* such an answer even if neither I nor anyone else as yet can give it, is how can there be self-awareness, our human style of self-awareness,

given that the being with such self-awareness can be understood, at least up to the point of self-awareness, through a resolutely materialist story?

So back to the story. I said we learned to have the self-awareness that we sometimes call 'self' or 'ego' or 'soul'. My story is that we learn to have such self-awareness using the same (or closely related) resources and the same (or closely related) inbuilt structures in our bodies that enable us to learn how to stick our tongues out. There's a wonderful experiment that was done in the 70's in which researchers demonstrated that neonates, if I recall less than two months old, just old enough for their eyes to focus on and track faces, can learn to stick their tongues out when an adult sticks out her or his tongue and holds the pose long enough. The infants struggle – I did this with my first child and it was amazing to watch his face wriggling, his mouth moving while he worked and worked at it – and then, suddenly they do it, they stick their tongue out and, as my son did, they giggle.

What's always stunned me about this experiment is this: it's amazing enough that the infant works at sticking their tongue out. What's really amazing is that the infant *knows* when they've gotten it right. When my son giggled the sense was overwhelmingly clear that he was so delighted to have gotten it right that he could only giggle. I giggled with him. But all the time I was wondering just how in the world he could possibly know he'd gotten it right.

I think the answer to that question is simply that our bodies are able to relate our experiences of ourselves to the experiences we have of other people and that our bodies do that with the very same sorts of entirely material neurophysiological processes that let us recognize faces, apples, trees, etc. When we see another person smile we can not only recognize their face as a face that's smiling, but we can smile as well and be confident that we're doing something that the other will see as the same as what they're doing. My story is that this is built into our neurological capacities, which develop through the same sorts of elaborate but purely material chains of chemical reactions that produce our abilities to recognize things in the world and successfully navigate the world populated with those things. Our bodies make it possible for us, when we see someone waive at us, to waive back in a way that answers the waive because we can feel what we'd be doing if we waived the way that person waived.

And it's not a big step from such capabilities, which I believe are all amenable to resolutely materialist explanations having to do with various neurological (and other, e.g. pheromonal response) structures we have (i.e. capabilities that grow within our bodies as we mature – by some blend of nature and nurture), to the capability to see ourselves like we see others seeing us and to have the confidence that the way we see ourselves is (at least more or less) like how others would see us could they see (experience) what we see (experience).

I know there's much more to be said to make this all sufficiently explicit to be tested with concrete, realistic experiments (and not things that rely on questionnaires), like a resolutely

materialist story should be. But I hope that the idea seems at least potentially plausible that our inborn biology is sufficiently rich and sophisticated that the sense that we can each be aware of each other, and each can be aware of ourselves as one of whom others can be aware, is something the spark for which is available very shortly after (if not literally concurrently with) birth.

This is the way, I should say one way, there can be a materialist explanation of consciousness that doesn't slight our continual sense of awareness. Consciousness is the light shone upon our being by the 'imagined' viewer of ourselves, the viewer who is ourself. The story I'm telling is that we each learn to look upon ourselves like we see others looking at us – 'looking at me' being no more mysterious than 'sticking my tongue out' is. But the result of me being one who looks at me is very different from me being one who sticks his tongue out, and it's that difference that makes the possibility of conscious awareness.

Full conscious awareness grows from that seed of being able to see ourselves (and through ourselves everything around us) as though we were another seeing us while still being ourselves. I want to underscore something: I am *not* proposing this as a 'solution' to 'the hard problem'. I *am* saying that for me, at least, this sort of story suggests how the conscious awareness we all value so highly might arise and exist as a physical phenomenon, a property of the right aggregate of matter in the right circumstances, and it has, moreover, some important characteristics that I think any successful story about consciousness has to have.

It has that effect for me because in the story of the infant's gradually emerging awareness of him- or herself as another of the sorts of beings among whom the infant exists, an awareness that arises from the infant's experience of being able to do what he or she sees others do, the infant's experience of recognizing success at such endeavors and the infant's experience of the others' reactions to his or her success as being another such thing the infant shares, in that story lies the opportunity for the infant to come to see himself or herself as seeing (and hearing and feeling and every other sensory modality) himself or herself just like others see himself or herself, and in that opportunity, which requires nothing but clearly material processes applied to different, but still material, circumstances from those with which such process may have engaged in evolutionarily prior animals, lies the possibility for the emergence of the light of awareness that each of us has.

Because once we're there, once we have the sense that we can each see each other as capable of doing the same and experiencing the same and that we can see that we each see the other that way (the baby smiles at the person with the protruding tongue and giggles just like the adult giggles and both feel the pleasure of a novel success), the rudiments of self-awareness, of the light of conscious awareness we each experience, are there.

Moreover, there is also present, it seems to me, the impulse to be as the others around us are. That impulse is manifest in the infant's attempt to stick its tongue out in response to the adult doing so. It's also manifest in the myriad other ways infants begin to imitate those around them. By calling it an 'impulse' I'm indicating that it's not something anyone deliberates about, just like we don't deliberate about breathing, just like an infant doesn't deliberate about sucking (nor we adults deliberate about chewing once the food is in our mouths). It's simply one of the factors that shapes our choices about what to do.

I call it an impulse to be as others are rather than an impulse to imitate because I think there's more to being as others are than just imitating what they do. 'Imitating' implies, to my ear at least, that there is some whole being that has an option in its wholeness to imitate *or not*. But what the infant is developing, what the child, adolescent and even adult continue developing is that being that can have the option to imitate. What the infant is doing, then, is something more than imitating. It is developing a way of being that is one manifestation of how humans can be human. That's too broad in certain respects. No infant develops into a universal human; we all develop into a human being that is an instance of the human beings among who we grow up.

I'm deliberately dodging the word 'culture' as well because, to my ear, it has a similar characteristic to 'imitate'. 'Culture' carries some at least potential implications of optionality, like 'imitate' does. We 'acculturate', as adults, to environments in which there's some meaningful sense in which we continue as the human we were before we

acculturated, and there's at least the potential for us to resist or even completely reject the forces of acculturation as an adult (we can, in the right circumstances, simply leave and go live our lives elsewhere). But the infant not only has no such independence, the infant can only become an adult human in the midst of other humans. Indeed an infant can only grow to have those capabilities that may let an adult strike out independently by developing them in the particular way that fits in the right respects with the humans among whom the infant grows. "Fits" and "right respects" here simply means in the ways that the interactions among the adults and infant prompt the infant to act as the adults around him or her expect him or her to act.

And in due course, the impulse to be as others are, the capacity to adapt our neurological repertoire to include the actions of others (aka imitation) and, crucially, the motivations others exhibit (and which recent experiments also suggest infants are very attentive to at a surprisingly early age) begins coalescing into language.

Language has some very peculiar properties that I don't notice getting much attention in the academic literature. First and most importantly, we have a sense that we know what words mean and that what we think they mean everyone else also thinks they mean. But how can this possibly be right? I mean I'm certainly assuming that about you and me as I write this (even though I don't know who you are as I type), and if you're still reading then almost certainly you are also assuming it about us. And I could say that long habituation to the usage of words and to how others respond when I use them right has solidified that

sense of sameness of understanding. OK, maybe that's right for me the 70+ year old adult, but what about when I was 1 or 2 years old and just starting to talk? Where did that idea of sameness of meaning come from?

And elaborating on that further, where did the idea of meaning come from – that's a complicated concept. For something to have a meaning there has to be both the meaningful something, let's call it W for short, and several relationships W has to the world, to other W's and to the speakers and hearers involved in the exchange. And if we're going to talk about all those relationships, we're going to want a short way to refer to them, call it 'meaning'. But how did the infant get from chortling and babbling to such a concept?

There are further problems about how it's possible to recognize individual words – even the simple ones infants mostly engage with like cat, doll, hat, ball, etc. – as the same. It's taken a lot of software work by hundreds if not thousands of researchers and programmers over the course of several decades to develop computer systems that can recognize speech, and even now that they can I don't believe anyone would argue that the way computer systems do it is the way our bodies do it. Be that as it may, the difficulty of attaining the current level of speech recognition is evidence for how *dissimilar* two people's pronunciations of the same word can be and how dissimilar even the same person's pronunciations on different occasions can be. So how does the infant do it?

I think these problems get far more tractable if we turn our perspective around and assume that what has to happen isn't that the infant learn about some discrete separate things called words and about how we say (tokens of) those words, but rather that the infant has to learn how to act like an adult. The mechanisms of speech are just muscular movements involving the air flowing through our mouths, noses, throats and into and out of our lungs, and in that way they're just like the tongue or smile. The infant learns to do what the adult does (and the same internal means of recognizing success apply here as well), so what's needed isn't a decoding of a sound stream (that sends us on a search for what's the same between two sounds) but rather a recognition of what the other *did* to make the sound – we hear the motion of muscles just like we see the motion of muscles in the smile or the protruding tongue.

And along with moving one's muscles the right way, there are a lot of other considerations that go into how to use language. Among them are what one is doing when talking and when it's appropriate (or not) to talk. I single these out because the fact that there are reasons to talk and times when it is or isn't appropriate to talk prompts the need to be able to talk about talking, and the need to talk about talking requires that we not have to talk about all the nuances of making the sounds just right (besides a lot of those nuances aren't available to our conscious attention let alone our deliberate decision-making capacities). It requires, in other words, some form of shorthand, a way to talk about talking without having to talk about everything else all at once.

Moreover, it requires that there be right and wrong ways to talk. I don't mean anything particularly grandiose about 'right and wrong' here, just that the infant has to learn "how it's done" and how it's done means there are ways it isn't to be done. Just what 'it' is that's done some way and not others has a very wide range – from pronunciation to tone of voice to timing relative to others' actions and innumerable nuances and conditions.

That there can be a right and wrong and that we need to have a shorthand to allow us to teach the infant (or each other) the right way leads, in my view, to the idea of words with meanings.

In my story, this process of learning language is an instance, certainly a crucial and perhaps even the central instance of how the impulse to be one among the others shapes the child's adaptation. The infant learns to be one who recognizes things that others recognize and that get called words, and the infant learns to be one who experiences those things as calling attention – the infant's attention and others' attention – to aspects of the world, with the expectation that there is some reason for directing their attention beyond simply the fact of the words having been spoken. Reference, in other words, is a way of orienting ourselves to one another and the world, and in that world one of the sorts of things to which we might call attention are the words we are speaking.

But one of the most important points about all this, in the context of a rigorously materialist understanding of being human, is that as infants we are learning to enact the existence of

fictional things. There's nothing in the world that pins down even a single instance of a word (a 'token' in philosophical jargon) other than our enacting its existence. No being other than human beings has any sense of a word as being what we all know it to be – a word. There's nothing in the world that delineates the word, nothing that marks two instances of a word as the same word, other than our human actions with respect to our own actions when speaking and our human actions with respect to the results of our actions, the marks on paper, when writing.

When I say 'enacting' what I am getting at is that we shape our actions in the way they would be shaped were there a real thing that is a word or its meaning. There's nothing else going on. And it's because we can recognize each other's actions (and motives and feelings) – including the residues of some of those actions (ink on paper) – as being like actions we might make (motives and feelings we can have) that it's possible for us to enact such things.

A talented mime can create the vivid impression in her audience that there is an object on stage with her even though everyone in the audience (and the mime, of course) can see that there is no such object. That's what words are, in my view. They're objects of a peculiar sort the existence of which is demonstrated by our actions involving those objects – and by absolutely nothing else (witness the difficulties with speech recognition systems).

So as we learn to see ourselves as others do, and as we learn to talk with them, we are learning to enact the existence of fictional beings in just the way that those we live with do it. Many of these beings have some form of physical presence – the spoken or written word – so their fictional aspects can't be just anything – but as we continue developing there need not be such a physical presence. The word isn't the spoken thing, it's something else and it 'has' a meaning (whatever 'has' means here). And our self, our ego, isn't our body, it's something else that our body 'has' or that has our body (whatever 'has' means here). I think that what we do as we mature, as we grow into having a fully specific self, is that we learn how to enact that self's existence without there having to be anything concretely real about that self apart from the body that 'has' it, and we learn how to enact that by doing as we see the others do whose seeing of us prompts the beginning of the process of enacting the existence of ourselves and seeing ourselves doing what we do.

There's a very important point about things that don't exist but the existence of which we enact. Such things don't have to be consistent, they don't have to stay the same, to any degree more than is required for us to continue experiencing their (continued, consistent, same) existence (where 'us' here means both each of us individually and the group of people whose actions are relevant to our perception of that continued existence). That qualification doesn't mean the thing can have any quality we want it to have; it means simply that there's no way to guarantee the consistent identity of such things. They can change without us noticing. The simplest example is the change in pronunciation of words that takes place over generations without people failing to be able to understand each

other. It still sounds like the same word but if you had a tape recording of Shakespeare pronouncing a word and a contemporary British person doing so the differences would be striking.

What this qualifier means is that while we each maintain a sense of the continuity of ourselves, and while there's something objective about my body being the sole successor of what I would have called my body 10, 20 or even 60 years ago (and in that sense still being the same body), nonetheless there's no constraint on our perception of ourselves being the same person other than whatever it is that we need in order to feel like we are still the same person. Some aspects of that are certainly the physical continuity that our living processes sustain, but others are to some large extent independent of that physical continuity. It's those other aspects whose actual continuity I'm saying has no constraint beyond whatever is necessary for us to retain the sense of our own selves continuing.

Pulling together several aspects of this – what we learn to do when we learn to talk is we learn to act in ways that display in some way, more or less specific, more or less detailed, our orientation to the world or some part of it. The individual structures of the discrete components (words, phrases, sentences) that comprise the language we each imagine we've acquired by adulthood and the interconnections among them are all aspects of that display. Our individual (and consequently collective) enactment of the existence (and meaningfulness) of those structures is what lets us display to one another how we are orienting ourselves to reality (including in that 'how' the display of which aspects of reality

we're orienting ourselves towards) which in turn displays if perhaps indirectly how reality is (or at least how it may be if our orientation is one our audience comes to share).

It's through those displays that we enact the existence of the things with which we do the displays (words, sentences, etc.) and it's through those displays that we can enact the existence of other things – institutions, ideas, theories, etc. – that also have no physical existence. I call them 'displays' because what I see us doing is manifesting ourselves in our actions. When we talk, those actions result in sounds and gestures, but actions don't stop there, of course. There are lots of actions we can take to display something – we can make marks on trees to show paths through the woods, we can stack stones to mark in an important location. All such displays are displays only for the others who share our sense of how it is one makes a display.

Displays themselves become the opportunity for action, we can develop our displays, say things more eloquently or more accurately. With the result that making displays can itself become the work to be done – call it art. Whatever else a work of art may be, I see it as a human display for humans. The physical object – that begins its existence in the deliberate act of an individual human and that develops through the continuing actions of that (or other) humans – manifests in physical form what was born from the mind of those initiating and developing humans. It is mind embodied in the material of the object, embodied through some period of time but not developing as its creator's (or creators') mind(s) continued developing beyond the completion of its creation.

Because our material minds, the minds that can create the work of art, the minds that can experience it, those material minds are never still. They are constantly changing and never repeat themselves. Notwithstanding that constant change, we interact in ways that help us preserve the sense that there is repetition, that things are stable, and we depend upon that preservation for our continued existence, especially today with the spectacularly elaborated technological processes by which we get everything from the most mundane of necessities to the most sophisticated medical interventions.

When I say the result of our actions can be mind embodied in the material of an object, I am not saying that some non-material reality has come to occupy a physical form. I am, instead, calling attention to the thought that we can only rarely fail to see when something was made by human hands. We see each other's actions as actions we might have undertaken ourselves, were we only competent enough and so inclined. That's how the object can be mind embodied.

But our continued existence also depends upon the reality of constant change. The reality is that nothing is the same twice and, in a fight, reality always wins – which is to say we ignore change, we try to constrain change, at our peril every bit as much as we refuse to preserve stability at our peril. Judging just how and when we should change something previously stable is one of those decisions of the second kind I was talking about above. We can't put together a science that will give us the answers to such questions because

science is by its nature something that aspires to stability – i.e. it aspires to ‘get it right’ which, if once ever attained, wouldn’t have to be revised (like, say, basic arithmetic).

The work of art we stand and look at is always both – it’s stable (at least paint on canvas generally is when shown in a gallery) and yet it manifests constant change. The physical actions of creating it leave their immobile physical traces, but those immobile traces display the motion, the change that created them, a display shaped by the physical realities of the materials with which the traces are formed.

This is a way that mind makes itself material and if we can see our minds as resolutely material we can see ourselves, our lives, as elaborate, richly, staggeringly complex ongoing dances that transform material realities into the material presence of mind. What a wonderful thing to be able to do.