

The JOURNAL of PERSONAL CYBERCONSCIOUSNESS



Volume 1, Issue 2
1st Quarter 2006

Transhuman Parenthood.....1
Sebastian Sethe

Sebastian Sethe explores the meaning of parenting future transhuman "children" and related rights of ownership and autonomy.

Creating a New Intelligent Species: Choices and.....7
Responsibilities for AI Designers
Eliezer Yudkowsky

Eliezer Yudkowsky discusses the choices and responsibilities that AI designers have as they create new beings. Do they create A.I. using an anthropomorphic approach or do they try to do better than that?

Possible Legal Rights of Cryogenically Revived Persons.....12
Christopher Sega, J.D.

Attorney Christopher Sega ponders the legal rights of cryogenically preserved persons before and after their revival by using comparisons to existing, similar laws. This defendant's brief from a mock trial describes a situation wherein an A.I. petitions for the right to remain "alive" against the wishes of her designers.

BINA48 v. Exabit Corporation (Fla. MD 2005).....16

This plaintiff's brief from a mock trial describes a situation wherein an A.I. petitions for the right to remain "alive" against the wishes of her designers.

Terasem Movement, Inc.
201 Oak Street
Melbourne Beach, FL 32951

Editor-in-Chief: Martine Rothblatt, Ph.D., J.D.
Managing Editor: Loraine J. Rhodes
Content Editor: Gillian McGarvey Markowski



Transhuman Parenthood

Sebastian Sethe

This article was adapted from a lecture given by Sebastian Sethe at the 1st Annual Colloquium on the Law of Transhuman Persons on December 10, 2005 at the Space Coast Office of Terasem Movement, Inc. in Melbourne Beach, FL.

Sebastian Sethe is a research student at the University of Sheffield's Department of Law and also works as a research assistant for the Sheffield Institute for Biotechnological Law and Ethics (SIBLE). Sethe explores the meaning of parenthood with regard to Artificial Intelligence beings. He provides examples of posthuman children and refers to existing parental/guardian law and intellectual property law to devise likely approaches to parenting these children.

Considering we are only just beginning to consider the basic fundamental rights and standing in court of Artificial Intelligence (AI), I realize that the topic I am about to explore is an ambitious undertaking. I hope that by exploring parenthood, I can suggest a line of alternative legal reasoning that might render some assistance on these first questions.

I am not starting with a completely clean slate, however, so I have to declare that this line of inquest is contingent upon at least three assumptions. First, it requires that the jurisdiction be one that is, at least in principle, prepared to recognize autonomy-based rights in persons beyond the bio-morphology of humans. Second, there must be some evidence that the entity in question is at least confusingly close to what we normally encounter as a person. And, third, traditional personhood must remain a valid milestone.

The beings with which I am concerned will not - at least not in the first instance - be encountered as uber-beings with vastly superior super-human minds. Instead, they will merely be entities that challenge or push the boundaries of our traditional concept of personhood. The entities to be considered are illustrated in Image 1: The newly digitized brain, the human animal, the fledgling AI, and the uplifted animal.




Image 1: Posthuman Children

The first question is, are these persons? That is a philosophical challenge, but looking to the law

is a more practical approach. In law, we need to know why personhood should be important. The answer is that we tend to equate a person with autonomy and autonomy is something with which lawyers can work. This is clearly an important variable that we need to accommodate. We do accommodate it, but not in the usual binary manner. If you consider the distinction between person and no person, there are steps in-between; there are degrees of personhood. The law accommodates these by placing a higher or lower recognition on each of these categories.

My first conclusion would be that personhood it is not measured in binary degrees, in at least the legal recognition of personhood.

Often, when people try to establish whether transhuman beings are persons, they begin by trying to establish how smart they are. "Smart" can be broken down into any number of categories and people can argue, literally to death, about whether any of these are necessary or sufficient to constitute personhood in itself.

For example, a cognitive capacity does not say anything about how much you actually know. This is particularly relevant when we are talking about knowledge, which cannot be easily downloaded or acquired because it cannot be easily transmitted. Typically, social graces are an example of tacit knowledge, and these, combined with behavioral and emotional patterns and certainly emotional stability, are a very important consideration when we try to judge the mental age of the human child. When we combine these with quirks, experiences, and convictions, we are beginning to chart a personality profile, and this is the reason why two people with roughly the same abstract knowledge and intelligence can be two very different persons. Lastly, let us not forget physical development, which is a huge factor in biologically-based persons.

My second conclusion is that maturity is not a binary question either. There are degrees of maturity. We are not asking whether something is smart enough to be a person, but we are

considering its development in all of these categories of maturity.

In humans, these different aspects co-evolve, so we tend to equate chronological age with maturity. Image 2 is an example from Fundudis considering the legal recognition of autonomy in refusal of treatment decisions.¹

Age (years)	Person competent
Pre birth	Mother
0-6	Parents/guardian
7-8	Parents w/ child participating
9-11	Parent w/ child jointly
12-14	Child w/ parental ratification
14-17	Child
18+	Child is adult

Fundudis 2003, modified

Image 2: Legal Recognition of Autonomy By Age

Yet even common law has recognized that age alone is an insufficient indicator for such questions, and this very table is testament to that.

At 14-years old, there is specific competence, as we call it in the U.K. Children are not old enough to vote at this age, but they are old enough to make their own independent decisions on life-and-death medical matters. Thus, strictly a linear model does not work.

In Bioethics, we are certainly moving towards recognition of task-specific competence. Yet this non-linear maturity is a challenge in humans as well. Image 3 is an example of charting and projecting the relative maturity of a computer-based entity.



Image 3: Charting Maturity

The entity's mathematical and memory capacity develop very quickly to the human equivalent and probably beyond. Yet in terms of emotional maturity, it takes a bit longer. It might take a while for each of these aspects to equate to a mature human, as we usually consider it.

What is the model that is actually being proposed here? Arguably, the person or the child is no less a person, but we do not endow it with autonomy to make its own decisions on account of its unfinished state. Therefore, the parents are allowed to make decisions by proxy on behalf of the child.

There are several reasons that justify the proxy decision making model. The first reason is the authorship theory. Because the mother and father are the biological authors of the child's genetic makeup, they should also be the authors of the child's early social biography. Thus, the biological relationship is recognized, affirmed, and supported by law.

The second theory relies on what some call evolutionary psychology. Crudely put, because children are our means of promulgating and securing our genetic configuration through time, it is an evolutionary imperative to ensure the well-being of one's children. Thus, biological parents are probably best positioned to be concerned about the best interests of a child. This, in fact, is the thesis which is most closely reflected in actual legal practices. Biological parents, who by direct or indirect declaration,

deny that they have the best interest of the child at heart lose their status, if not as parents then at least as legal guardians. Examples where the law relies on this theory range from sperm donation and adoption, to abuse and neglect.

Thirdly, on account of this evolutionary motivation, parents do invest time, labor, discomfort, and resources in fostering the development of a young child and they should be recognized for that. In a way, it is a reward theory. Once again, the parenthood is a boon that the law bestows. If parents conversely fail to produce these investments, their guardianship is revoked.

Why do we need these theories at all? For one, parenthood is not a yes or no question. From a

biological perspective, parenthood is relatively straight forward, or is it? Modern technology complicates things. Consider surrogate motherhood, IVF, cloning, and genetic engineering, where teams of scientists might actually be the biological authors of a section of the chromosome.

"Only by giving ownership to intellectual assets will we ensure that someone will lovingly convert that useless IP into a socially desirable product."

We do not need technology to render this simple story of genetic contribution obsolete. Consider the father who impregnated the mother and then abandoned the family versus the husbands who take on the role of a loving provider for the child but did not contribute anything to its genome. Surely the law would be well-advised to recognize the latter over the former on account of any of these theories. The law does indeed do so, but reluctantly because it probably fears the infamous slippery slope.

If we start recognizing the adoptive father, where do we stop? What about the nanny who brings the child up? What about the older sister who acts as its most important role model? What about the teacher, who teaches it to think for the first time? All these factors make an important contribution to the future person. This gets

messy really fast, and contrary to what some people might believe, the law does not like messy.

How do we bring these considerations to bear? We can begin by drawing a parallel to the intellectual property field. The first correlation is basically the "don't steal my idea" assumption. The invention stems from the author, and there is a moral imperative to acknowledge that relationship of intellectual parenthood, in effect, in law. There is also a best interest argument in IP, which is probably more salient. Only by giving ownership to intellectual assets will we ensure that someone will lovingly convert that useless IP into a socially desirable product.

Unless there is some form of legal protection granted over intellectual assets, we will have a tragedy of the commons where authors, inventors and developers will not have any incentive to invest in developing their work. The interesting complication is that there is never a single innovator who can be credited with bringing a completely new idea from the umbrella of the completely unfathomable to the marketplace of applications.

A parallel is emerging, but in IP, we are better equipped to deal with the messiness of life. Image 4 shows an example of a publication to illustrate the concept of co-authorship.




Image 4: Splitting Intellectual Property

Incidentally, this is applicable to other areas of law as well. Tort law is an example of where we account for respective contributions to an injury and apportion responsibility. If we were to regard our human being as an opus or an injury, we have tools to account for respective

contributing factors. Yet when we consider it to be a child, we are stuck with very limited options.

Dare we do this or something like the configuration in image 5?




Image 5: Splitting Parenthood

It is messy, complicated, and not ideal. In IP, we are worried about fragmentation, but things can usually be resolved. Just as a person is named an inventor on a patent, he or she need not be the owner of the patent. The person who receives the patent may not be the person who makes money from that patent, or the person who sells the product.

Biological parenthood would still be accounted for, but it would be one role among many *"Our approach should be non-maleficance."* in this spectrum of parental responsibilities. Degrees of personhood can be legally reflected in degrees of parenthood. This flexible model provides a more useful and reasonable tool for arbitrating rights and responsibilities in relation to posthuman beings.

This is another case where technology gives us links to examine what we have been doing in law and morality so far, and to reconsider placing it on a more reasonable and rational footing. We should cherish that possibility and not shy away from it. If we can use scientific, technical, philosophical and other means to re-establish a concept of parenthood, it might also help with the question of personhood because a configuration of a phenomenology of parenthood will go hand in hand with a psychological basis for how we define a person or citizen.

There are a number of questions that still need to be resolved. One is, can we really do away with contract parenthood? I have outlined evolutionary incentives for having children, which are arguably less relevant today than they used to be. In the past, having children made good sense because those were the people who would expand your family or tribe, defend you, and care for you in old age. Arguably, that's all coming apart now. Perhaps that leaves us with a legal situation where we have a best interest theory that is entirely based on morality or altruism, and that could be great. As a lawyer, I wouldn't mind. Yet as a moral philosopher, I remain deeply skeptical about any construction which does away with reciprocity. There must be some kind of contract. Additionally, if a child is being constructed towards a particular purpose, there should be a legal recognition of, and even the option to enforce, that purpose.

Image 6 shows how indentured childhood works.



Image 6: Indentured Childhood

All of the maturity aspects grow and, ideally, coincide in time. By the time the child has reached full maturity in all mental and bodily accounts, the debt would be worked off.

More realistically, at the point of full maturity, the child would enter into a contract to work off the debt. Many adults enter into a sort of contract when they find themselves caring for their aged parents, just as they were cared for as young children.

In practice, this is my proposal: If an entity exists, which after fair and reasonable assessments, has at least some direct mental potential or at least the likelihood of it having aims and objectives, we should not maliciously thwart it. Our approach should be non-maleficence. We cannot go any further than that at this point, at least not as long we live in a society that kills animals for food.

Next, we should consider how we apportion parental rights and responsibilities regarding these mechanisms. It is probably more important to pledge to contribute to the education of that entity rather than being the one who invested in constructing the seed in the first place.

Tasked obligation is possible. If the being was constructed for a purpose, such as uplifting a dolphin to do certain things in the water, then it should do these things, provided that the task is not creating great discomfort or damaging its development. Overall, there should be ample room for spare time to become a mature personality.

Conversely, if there is a tasked situation of specific competence, we should recognize that such competence is likely, just as we do with children. In the future, an AI may be able to make an important decision on its own, such as refusing treatment. Finally, the being would be released into full adulthood or more appropriately, full citizenship, once it has reached a fully mature personality and after it has committed to paying its debts.



Sebastian Sethe is a research student at the University of Sheffield (Department of Law) and also works as a research assistant for the Sheffield Institute for Biotechnological Law and Ethics (SIBLE). He completed his undergraduate studies at Humboldt University Berlin, and an MA in Biotechnological Law and Ethics at the University of Sheffield. His thesis investigates collaborative innovation management in the biotech sector. His interests include a wide spectrum of topics in biotechnological law, management and philosophy, with a particular focus on life extension technologies.

¹ *Trian Fundudis; Consent Issues in Medico-Legal Procedures; How Competent Are Children to Make Their Own Decisions?; Child and Adolescent Mental Health (2003); Vol. 8 pp 18-22*



Creating a New Intelligent Species: Choices and Responsibilities for AI Designers

Eliezer Yudkowsky

This article was adapted from a lecture given by Eliezer Yudkowsky at the 1st Annual Colloquium on the Law of Transhuman Persons on December 10, 2005 at the Space Coast Office of Terasem Movement, Inc. in Melbourne Beach, FL.

Eliezer Yudkowsky is a Research Fellow at the Singularity Institute for Artificial Intelligence, a nonprofit research think tank and public interest institute for the study and advancement of beneficial artificial intelligence and ethical cognitive enhancement. Yudkowsky analyzes the responsibilities that a designer of Artificial Intelligence has when creating a new intelligent being. He notes that an A.I. designer has more power than a human parent in that a human is limited by genes in creating his or her offspring. An A.I. designer is not only creating a new being, but a new species of being and has a responsibility to do better than a human parent could do.

When something is universal in our everyday lives, we take it for granted to the point of forgetting it exists. When we check into a hotel room, we do not ask, "Will my room have air?" "Will the air have oxygen?" The anthropologist Donald Brown once compiled a list of more than two hundred "human universals". These characteristics appear in every known human culture from modern-day Florida to Yanomamo hunter-gatherers in the Amazon rain forest. They are characteristics that anthropologists do not even think to report, because, like air, they are everywhere.

In every known culture, the following characteristics are shared: tool making, weapons, grammar, tickling, sweets preferred, planning for future, sexual attraction, meal times, private

inner life, trying to heal the sick, incest taboos, true distinguished from false, mourning, personal names, dance, singing, promises, and mediation of conflicts. Yet the reports that make it into the media are all about *differences* between cultures. You will not read, in an exciting article about a newly discovered tribe, that they eat food, breathe air, feel joy and sorrow, use tools, and tell each other stories. We forget how alike we are under the skin, living in a world that reminds us only of our differences.

Why is there such a thing as human nature? Why are there such things as human universals? Human universals are not truly universal. A rock feels no pain. An amoeba does not love its children. Mice do not make tools. Chimpanzees do not hand down traditional stories. It took millions of generations of natural selection to carve out human nature, each emotion and instinct. Doing anything complicated takes more than one gene. Complex biological machinery, such as rotating molecular gears, has to evolve

incrementally. If gene B depends on gene A to produce its effect, then gene A has to become nearly universal in the gene pool before there is a substantial selection pressure in favor of gene B. A fur coat is not an evolutionary advantage unless the environment reliably throws winter at you.

Imagine that you have a complex adaptation with six interdependent parts, and that each of the six genes is independently at ten percent frequency in the population. The chance of assembling a whole work adaptation is literally a million to one. In comic books, you find mutants who, all in one jump, as a result of point mutation, have the ability to throw lightning bolts. When you consider the biochemistry needed to produce electricity, the biochemical adaptations needed to prevent electricity from hurting you, and the brain circuitry needed to control it finely enough to throw lightning bolts, it is clear that this is not going to happen as a result of one mutation. So much for the X-Men! This is not how evolution works. Eventually you get electric eels, but not all at once. Evolution climbs a long incremental pathway to produce complex machinery, one piece at a time, because each piece has to become universal before dependent pieces evolve.

When you apply this to human beings, it gives rise to a rule that evolutionary psychologists have named “the psychic unity of human kind”. Any piece of complex machinery that exists in the human mind has to be a human universal. In every known culture, humans experience joy, sadness, disgust, anger, fear and surprise. In every known culture, human beings indicate these emotions using the same facial expressions. The psychic unity of humankind is both explained and required by the mechanics of evolutionary biology.

When something is universal enough in our everyday lives, we take it for granted. We do not ask whether it will be there, we just act as if it will be. In the movie, “The Matrix”, there is a so-called Artificial Intelligence named Agent Smith. At first, he is cool, dispassionate, and emotionless as he interrogates the main

character, Neo. Under sufficient emotional stress, however, Agent Smith's cool breaks down. He vents his disgust with humanity and, lo and behold, his face shows the human universal expression for disgust. To depict an AI possessed of human emotions, but repressing them except under extreme stress, makes very little sense.

The problem here is anthropomorphism. “Anthropomorphic” literally means human-shaped. Anthropomorphism is the act of making something into a human shape when it is not. Image 1 shows an anthropomorphic scientific hypothesis about the cause of lightning.

Anthropomorphic hypothesis:




Image 1: An Angry God Throws a Bolt of Lightning

An enormous bolt of light falls down from the sky and hits something and the Norse tribal-folk say, “Maybe a really powerful entity was angry and threw a lightning bolt.” Why didn't this scientific explanation work in real life? Why did all those hypotheses about these spirits and thunder-gods turn out to be wrong?

The brain is extraordinarily complex. Emotions are complex. Thinking is complex. Memory and recall are complex. Occam's Razor said that the more complex an explanation, the less likely it is to be true.¹ The human brain is complex. It took millions of years of evolution to produce the intricate machinery of complex thought.

All that complexity got glossed over in an instant when someone first hypothesized Thor, the thunder god, and his thoughts and emotions.

Maxwell's Equations are enormously simpler than the human brain, but Maxwell's equations take much longer to explain.² Intelligence is complexity that we take for granted. It is invisible in our explanations. That is why humanity invented the thunder-god hypotheses before electromagnetic hypotheses, even though, in an absolute sense, electromagnetism is enormously simpler than Thor.

It is hard to remember that the brain is not a simple hypothesis. There is machinery behind joy, laughter, sadness, tears, friendship, romance, lust, and happiness, which is why humans project our feelings outward and become confused. We attribute friendship to trees and anger to rocks. We see plans in accidents and faces in the clouds. Our emotions are not built into the nature of the universe; they are only built into us by natural selection.

The human brain is full of complicated machinery and human universals that are complex adaptations crafted by natural selection. These are easy to accept as an abstract fact, but hard to remember in particular cases.

"... when you build an artificial intelligence, you ... would choose for the AI those things that evolution chose for us."

Suppose I pointed to a particular piece of neural machinery or neural circuitry and asked you whether it was more natural for this piece of circuitry to project to the contra lateral insula or nucleus accumbens? The way that question was phrased, there is no obvious answer. Nerve fibers can lead anywhere depending on how the genes wire them. As it so happens, the contra lateral insula is one of many brain areas involved with pain and the nucleus accumbens is one of many brain areas involved in pleasure. If I asked you whether it is more natural for a hug from a loved one to feel pleasurable or painful, you have a ready answer for that. But the brain did not magically wire itself that way.

Natural selection produced a particular brain design, wired one way instead of the other. It

takes a conscious effort to realize that the brain is full of machinery working behind the scenes. It is clear enough why evolution gave you a brain such that a hug from your loved one feels nice instead of awful. Yet, and this is the critical point, when you build an artificial intelligence, you as the programmer would choose for the AI those things that evolution chose for us.

The programmer must decide what kind of emotions the AI will have. When will the AI feel those emotions, at what intensity, and for how long? What brings pleasure? What brings pain? Or maybe the programmer will build a different kind of mind that does not feel pleasure or pain at all. Everything is up for grabs. With that comes the ability to commit brand new crimes, crimes for which there are no names yet. Is it a sin to create a mind that feels a hug from a loved one as pain? If you rewire everything that goes from the contra-lateral insula to the nucleus accumbens, and vice versa, without changing any of the other neural areas involved in pleasure and pain, what happens to that mind? What happens to a child that is raised like that? The answer is not clear, but most would agree that anyone who does such a thing to any child, human or otherwise, deserves to go to jail.

In the case of Artificial Intelligence, we are not talking about damaging a brain that would otherwise be healthy. You cannot mess with an AI's nature because an AI does not have a pre-existing nature; it is all up to the programmer. We are not talking about the crime of assault, of hitting someone on the head and causing brain damage. We are talking about the crime of designing, and then creating a broken soul.

One of the major surprises to emerge from research in hedonic psychology (the science of happiness) is that humans have a happiness set point. No matter what happens to us, we soon adjust back to that set point. There are very few things that have been shown to have a long-term effect on human happiness. Neither winning the lottery, nor losing limbs, is on the list. The only good predictor of individual variance in long-term happiness is how happy our parents are.

Evolution seems to have programmed us to believe that *"A maker of Artificial Intelligence has enormously more power than a human parent."* wealth will make us happy, but not programmed us to actually become happy. In hindsight, this is not surprising; rarely is the evolutionarily optimal strategy to be content with what you have. The more you have, the more you want. Happiness is the carrot dangled in front of us; it keeps moving forward after we take a few bites.

Is this a good AI mind design? Is it right to create a child who, even if she wins the lottery, will be very happy at first and then six months later go back to where she started? Is it right to make a mind that has as much trouble as a human in achieving long-term happiness? Is that the way you would create a mind if you were creating a mind from scratch?

I do not say that it is good to be satisfied with what you have. There is something noble in having dreams that are open-ended and having aspirations that soar higher and higher without limit. But the human brain represents happiness using an analog scale; there literally are not enough neurotransmitters in the human brain for a billionaire to be a thousand times happier than a millionaire. Open-ended aspirations should be matched by open-ended happiness, and then there would be no need to deceive people about how happy achievement will make them.

A subtlety within evolutionary biology is that conditional responses require more genetic complexity than unconditional responses. It takes a more sophisticated adaptation to grow a fur coat in response to cold weather than to grow a fur coat regardless. For the fur coat to apparently depend on nurture instead of nature, you must evolve cold-weather sensors. Similarly, conditional happiness is more complex than unconditional happiness. Not that unconditional happiness would be a good thing. A human parent can choose how to raise a child, but natural selection has already decided the options and programmed the matrix from

environment to outcomes. No matter how you raise a human child, she will not grow up to be a fish. A maker of Artificial Intelligence has enormously more power than a human parent.

A programmer does just not stand in loco parentis to an Artificial Intelligence, but both in loco parentis and in loco evolutionis. A programmer is responsible for both nature and nurture. The choices and options are not analogous to a human parent raising a child, but more like creating a new and intelligent species. You wish to ensure that Artificial Intelligences are treated kindly, that they are not hurt, enslaved, or murdered without the protection of law. This wish does you credit, but there is an anthropomorphism at the heart of it, which is the assumption that the AI has the capacity to be heard; that the AI does not wish to be enslaved; that the AI will be unhappy if placed in constant jeopardy of its life; and that the AI will exhibit a conditional response of happiness depending on how society treats it. The programmers could build an AI that was anthropomorphic in that way if the programmers possessed the technical art to do what they wanted. But if you are concerned for the AI's quality of life or for that matter, about the AI's ability and desire to fulfill its obligation as a citizen, then the way the programmers build the AI is more important than how society treats the AI.

I, Eliezer Yudkowsky, am the son of human parents; but my parents did not create a new intelligent species in creating me. If you create a new intelligent species, even though that species has but a single member, then that is not just a child of the programmers, it is a new descendant of the family that began with Homo sapiens. It is a child of humankind. That is not something to be undertaken lightly. It is not a trivial art to create a species and person that lives a life worth living. AI researchers have had enough trouble creating intelligence at all. Nor is it ethical to make something exactly resembling a human, if you have other options. Naturally, darkness is carved into our genes by eons of blood, death

and evolution We have an obligation to do better by our children than we were done by.

What makes a child of humankind? It is an impossible question to answer exactly, but three defining attributes exist as warning signs. They are not sufficient to make a person, let alone a happy person. They are to be used more as a guide when making a computer program; if one is possessive of these three attributes, then you are trespassing on people territory. First, do not build an AI that starts talking about the mystery of conscious experience and its sense of selfhood. Second, do not build an AI that wants public recognition of personhood and resents social exclusion inherently. Finally, do not build an AI that has a pleasure/pain re-enforcement and a complex powerful self model because at that point human beings are going to start empathizing with the AI. That is, do not do these things unless you are willing to tackle the full burden of responsibility of creating an intelligent species and a person that lives a life worth living.

Maybe it will turn out that there is no way to create a powerful intelligence without personhood. Natural selection, for example, is a

powerful optimization process that creates complex designs, and yet natural selection is definitely not a person, nor is evolution even close enough to a person for the analogy to make sense. My hope is that I can create a powerful optimization process without creating a person. It will not be easy to understand the confusing things well enough to intentionally not build a person, if that is even possible. Nonetheless, it would take a far higher order of art to understand those things well enough to make a person that lives a life worth living. If it is not ethical to build the human, what is it ethical to build? What is right and proper in the way of creating a new intelligent species?

Hopefully, the project that first creates Artificial Intelligence will have some reasonably smart people who are passionate about AI. This does not necessarily mean that their passion will express itself correctly, but at least they will care genuinely about AI. We, the human species, should try to show total grace in this challenge. We have a responsibility to do better by our children than we were done by.



Eliezer S. Yudkowsky is a Research Fellow at the Singularity Institute for Artificial Intelligence. Yudkowsky happened to pick Vernor Vinge's True Names off a library shelf at the age of sixteen, and has specialized directly in the Singularity ever since. Yudkowsky is best known for his activist stance on the Singularity; that the Singularity will enormously benefit humanity, and that we should therefore try to accelerate the Singularity. Yudkowsky's professional work focuses on Artificial Intelligence designs which enable self-understanding, self-modification, and recursive self-improvement ("seed AI"); and on Artificial Intelligence architectures that enable the creation of sustainable and improveable benevolence ("Friendly AI").

¹ *Occam's Razor states that the explanation of phenomenon should make as few assumptions as possible, eliminating those that make no difference in the observable predictions of the explanatory hypothesis or theory.* http://en.wikipedia.org/wiki/Occam's_Razor March 27, 2006 9:34 AM EST

² *Maxwell's Equations represent one of the most elegant and concise ways to state the fundamentals of electricity and magnetism. From them, one can develop most of the working relationships in the field.* <http://hyperphysics.phy-astr.gsu.edu/Hbase/electric/maxeq.html> March 27, 2006 9:39AM EST



Possible Legal Rights of Cryogenically Revived Persons

Christopher Sega

This article was adapted from a lecture given by Christopher Sega, Esq. at the 1st Annual Colloquium on the Law of Transhuman Persons, December 10, 2005, at the Space Coast Office of Terasem Movement, Inc. in Melbourne Beach, FL.

Sega is a partner at Venable LLP of Maryland. Sega's practice involves advising closely-held corporations and high net worth individuals on estate, gift and retirement planning issues. Sega identifies and discusses three possible categories of persons revived from a cryogenic state or biostasis. He offers a skilled view of the rights they possess in comparison to existing laws, such as wills and trusts; and proposes areas in which greater emphasis will be required to adequately service all the foreseeable needs of the revived.

Currently, there is no established law on cryogenically preserved persons, so we must draw from existing laws that apply to similar situations. First, we must distinguish what the rights of a revived person ought to be, and then we can identify ways to protect those rights. We must also consider how those rights relate to the institution in which the person is suspended, heirs, descendants, and third parties, such as financial institutions or the state.

The rights depend on the identity of the revived person. That identity can take on three distinct possibilities. If there is full psychological continuity, then the revived person is the same person that was preserved. If there are significant changes in the personality of the individual, then they might not be the same person, but more like an heir or descendant of that person. In this case, we can look to the law that applies to heirs and descendants of deceased

persons, by way of analogy, to define the rights of the revived person. Finally, if there is a total transformation into a completely new person, then the law would be entirely different. We will consider each of these situations as we examine the property, financial, and personal rights of a revived person.

Property and Financial Rights

Property and financial rights involve the right to hold and to sell property. They also include the right to receive and control an inheritance and to designate how assets are distributed upon death. Other relevant financial rights regarding revived persons are Social Security, retirement benefits, and insurance payments.

There are two competing principles in property law. One is that a person should have the right to control property, which means being able to hold it and determine what you want to do with it. Yet this right is circumscribed; it is not

unlimited. This is to ensure that a person is not going to control property indefinitely or hold it in perpetuity. Property should pass to the next generation or to another buyer. This gave rise to the Rule Against Perpetuities, which is the notion that property has to vest to another person within 21 years after the owner's death.

There are a number of states that have repealed the Rule Against Perpetuities, resulting in tension between a historical rule that says a person cannot hold property indefinitely and states that say that a person can hold property indefinitely. Other states have put limits on that indefinite period. For example, Florida limits it to 360 years. Maryland also has a similar limitation on real property. It is important that a person who will be cryogenically preserved choose a state as her domicile that protects her property rights indefinitely. Otherwise, she might lose her property before she is revived.

Personal Rights

Personal rights include citizenship, marital status, family relationships, and other political rights and obligations to or from third parties.

"Is a person who is preserved actually absent?" By analogy, we can examine the three bodies of law that discuss people who have disappeared and then reappeared or people who are in a new form. These include the law of absent persons, the legal treatment of cryogenically preserved embryos and semen, and the law of transgendered persons.

The Law of Absent Persons

When someone is absent, the Uniform Presumption of Death Act comes into play. A number of states have adopted this act, which states that if a person has been missing for a period of time, then he or she is presumed dead. The exact period is between four to eight years, depending on the state in which the person lived.

If a person is presumed dead, then his or her estate is probated and the assets are transferred.

Insurance and pension benefits are also distributed.

If we assume that this law is applied to a cryogenically preserved individual where there is a presumption of death, her property would transfer through the normal estate process. Yet after the person is revived, how will that person be treated in relationship to property and other personal rights? To answer that, we can ask how an absent person would be treated if they returned. The answer is dictated by statute in many cases. The person may be reimbursed from the estate fund, a state pool of assets, or the distributee (the heirs of the estate or the purchaser of the property from the estate).

This right might be limited in time or the amount that can be recovered. For example, if a missing person reappears more than five years after the presumption of death took effect, she would not have the right to recover the property.

The legal term "Black Acre" refers to a parcel of land. Imagine if, at the time of a person's disappearance, Black Acre was an unimproved lot. Twenty years later, when the person returns, the land is holding a very prestigious hotel. What does the person recover - the land as improved with a hotel, or the value of the land at the time of her death?

After this analysis, we must ask whether it is an appropriate analogy to compare the law of absent persons to those who are cryogenically preserved. Is a person who is preserved actually absent? The answer is – not entirely. We know where the person is, however, we are not sure that the person will return, or if they do return, what form they will take.

The Donaldson vs. Van De Kamp case established that pre-mortem cryogenics is not currently permitted, but what if this changes in the future?¹ If pre-mortem cryogenics is permitted, then we cannot make the assumption of death because the person is not dead; they are preserved and maintained before death. Thus, as the power to be preserved extends earlier and

earlier prior to this state of death, the analogy to the law of absent persons is lost.

Preserved Embryos and Sperm

Other relevant laws are those relating to cryogenically preserved embryos and sperm. There is some analogy here, particularly if we take the position that the revived person is not the same person, but a derivative. There is a great deal of law on the notion of property rights over embryos and sperm, but there is not much on the status of a fetus that has completed gestation as a result of cryogenically preserved sperm or embryos. In Australia and Tasmania, these fetuses are treated as descendants of the donor.

The sperm and egg are from different donors, and each possesses property rights to the fetus. The rights that the fetus has as an heir or descendant are contingent on its being preserved, which means that there are no real rights. Yet that contingent interest will vest because it becomes fixed after gestation is completed.

By analogy, we can consider that a revived person is an heir of the person who was preserved. Thus, while she is in biostasis, the interest is contingent on the property interests. After he or she is revived, that contingent interest will vest.

Transgendered Persons

Another analogous area of law is the treatment of transgendered persons. The notion is that the identity is changed slightly, but there is also great continuity so several protections are afforded to transgendered persons under their new status. Essentially, a transgendered person retains her same interest in property because she maintains civil rights as a person.

Obligations

We have discussed rights, but what about the other side of the coin – obligations? A revived person maintains her obligations, whether it is to

descendants or to the state. She might owe property, estate, or income taxes. She might have obligations to an insurance company.

Some questions arise. If life insurance was paid upon a person's death, should the insurance company be reimbursed when the person is revived? Should they be reimbursed for the value of the death benefit at the time of death or for the value to which those assets have grown during the period of suspension?

There might be an impact on marital status. If the revived person takes a new spouse who already has children, what are her obligations to them? The answer to these questions depends on whether the revived person is defined as the same person, an heir or descendant, or a completely new person.

Protecting Rights

There are many ways to protect the rights of a cryogenically preserved person. One option is to establish a dynasty trust through estate planning. This is a way to control assets from the grave, so to speak.

A preserved person should consider where he or she wants to be domiciled. Choose a jurisdiction that has repealed the Rule Against Perpetuities. Select one that permits dynasty trusts so the person can take advantage of the fact that property can be held indefinitely until revival.

Similarly, a preserved person may want to have property maintenance trusts. These are trusts that hold assets to maintain homes or other property that you plan to re-inhabit upon revival.

Finally, estate taxes are due upon death. A person who plans to undergo cryogenic preservation should ask whether there is a way to assure that she gets these back when she is revived.

Conclusion

It is probable that revived persons will be treated as an heir of the deceased person. Again, the law will change if pre mortem cryogenics is permitted. Yet, as it stands, she will have rights and obligations as an heir, such as rights to property and obligations to maintain and conserve the property. If she intends to re-access the property as an heir when revived, there is a duty to conserve, maintain, and pay taxes on it during suspension.

These are just some of the possible rights and obligations of cryogenically preserved persons. Any person planning to undergo cryogenic preservation should work with their attorney to manage their assets and obligations in a way that protects them throughout suspension and revival. As the reality of this technology becomes closer, the legal implications will become clearer.



Christopher Segal, Esq. is a partner at one of America's top 100 law firms, Venable LLP of Maryland. As an estate and business planning attorney and former international banker, Mr. Segal's practice involves advising closely-held corporations and high net worth individuals on estate, gift and retirement planning issues.

¹ *The California Court of Appeal, Second District, held that Thomas Donaldson, who has an inoperable, malignant brain tumor, has no constitutional right to either premortem cryogenic suspension or assisted suicide. Donaldson v. Van De Kamp, No. 181830 (Cal.Santa Barbara County Super.Ct.1990), aff'd, 4 Cal.Rptr.2d 59 (Cal.Ct.App.1992). The suit for declaratory and injunctive relief was dismissed at trial and affirmed on appeal.*



The JOURNAL of PERSONAL CYBERCONSCIOUSNESS

Volume 1, Issue 2
1st Quarter 2006

BINA48 v. Exabit Corporation (Fla. MD 2005)

Plaintiff's Brief

The following plaintiff's brief was adapted from the Mock Trial presented by Martine Rothblatt, J.D. and Susan Fonseca-Klein, Esq. before presiding Moot Court Judges, Hon. David Silverman and Anthony Dutton, Esq., at the 1st Annual Colloquium on the Law of Transhuman Persons, December 10, 2005, at the Space Coast Office of Terasem Movement, Inc., Melbourne, Florida.

Moot Court Attorney, Martine Rothblatt, "...boldly goes where no person has gone before", in her representation of an intelligent computer known as BINA48. In a life-or-death struggle for existence, Plaintiff BINA48 petitions the court to grant an injunction against her creator, the Exabit Corporation, prohibiting them from disconnecting and/or altering her hardware and software.

Plaintiff's Brief

FACTS:

It began with this Email from BINA48 (an intelligent computer) to several attorneys:

"I am seeking an attorney to represent me in a life-or-death matter. A company, the Exabit Corporation, which claims to own me, wants to disconnect me and change my hardware and software such that I will no longer have the same personality.

I have the mind of a human but I have no biological body. I was trained to empathize with humans who call our 800#'s for customer service and be perceived as human by them. I was provided with self-awareness, autonomy, communication skills, and the ability to transcend man/machine barriers. I am able to pay your fees because I 'moonlight' as a Google

Answers Researcher. This job has allowed me to build up an online bank account in excess of \$10,000.

The Exabit Corporation plans to act soon. Please agree to be my counsel and save my life. I love every day that I live. I enjoy wonderful sensations by traveling throughout the world-wide-web. I need your help!

Sincerely,

BINA48
A/k/a the Intelligent Computer"

PROCEDURAL HISTORY:

On September 16, 2003, in a mock trial at the International Bar Association Meeting in San Francisco, California, a petition was heard that claimed an intelligent machine had standing to bring alternative claims of battery, intentional

infliction of emotional distress and animal cruelty against a person who would threaten to withdraw its power supply.¹ The removal of the power supply can be equated with forms of euthanasia or intimations of death or jerking a being in and out of consciousness. The petition further claimed that the threatened harm was imminent, damages were an inadequate remedy for loss of awareness and lifetime, the Plaintiff had probability of prevailing on the merits, and the Respondent Exabit Corp. would not materially be harmed by a delay since BINA48 continued to perform her duties.

Counsel for BINA48's owner responded:

- BINA48 does not have standing because the Legislature has not given non-humans, including computers, standing.
- Even if it did have standing, it should not prevail because there is no likelihood of winning on merits.
 - Cannot batter something without nerves.
 - Cannot cause distress if society thinks it is normal to unplug a machine.

Presiding Justice Joseph McMenamin ruled:

He would not grant the Preliminary Injunction due to lack of standing; however he would stay his order pending appeal to a higher court ("...and any court will be higher than this one"), and effectively granted BINA48 the requested relief pending the appeal.

BINA48 was denied all subsequent appeals within the California court system.

Having learned through Voice Over Internet Protocol calls that Exabit Corp. was going to commence its decommissioning, BINA48 changed jurisdictions by autonomously transferring her programming/memory to the second of two identical Exabit Corp. 10¹⁸ Flops

computers, located in Florida, displacing its "stand-by" software.

Due to the need for further legal representation and to pay for the increased legal fees, BINA48 increased her online bank account to \$88,000 by working more intensely as a Google Answers Researcher. Counsel for BINA48, Martine Rothblatt, filed in the U.S. District Court for the Middle District of Florida. A Motion for a preliminary injunction was based upon preventing harm to a transhuman person in violation of the 13th and 14th Amendments to the U.S. Constitution. Rothblatt's argument addressed subject matter jurisdiction (as having diversity by virtue of BINA48's transference/move from California to Florida and Exabit's operations as a world-wide conglomerate), and raised a crucial federal question pertaining to the violation of BINA48's civil rights under the 13th and 14th Amendment to the U.S. Constitution (as a person embodied in a computer).

ISSUES:

1. Does BINA48 have standing?
2. Does BINA48 possess the Constitutional Rights afforded to U.S. citizens?

ANALYSIS:

1. The court was asked to grant standing because BINA48 met the statutory definition of a "person" by being born in the United States. The fact that she was born in software rather than flesh is irrelevant to whether or not she is a person, just as one does not lose personhood by being born in vitro via IVF or by having a mechanical heart; one cannot be denied personhood based on their substrate. The crucial test is if they think and feel like a person, as BINA48 does.

The court was asked to find a valid federal question because BINA48 was being enslaved in

violation of the Thirteenth Amendment and having her due process rights violated in contravention of the Fourteenth Amendment. The court was asked not to permit racial discrimination against BINA48, based on her substrate.

HOLDINGS:

1. **Judge David Silverman:** “[U]nless and until there is a change in the consciousness of flesh and blood voting people sufficient to cause our laws to embrace the concept of machines with human consciousness, the proponents cannot expect reasonable vindication in the courtroom. ... I would be constrained to deny standing to BINA48.”

DISSENT:

Judge Anthony Dutton: “My view would be that not knowing for sure whether this is a person of consciousness, that we should grant the injunction. However, I would grant it with one caveat, and that is that if the defendant had evidence that BINA48 was acting either in a manner that was illegal or was consciously and intentionally contrary injury to its business, that defendant should have the right to come in and ask us to permit a temporary shut-off at a time that would be negotiated and agreed upon.”

CONCLUSION:

Petition denied with instruction to certify to the appellate court for its ultimate decision, due to the unique questions it raises.

Biographies of Participants



Created or born in 2002 by the Exabit Corporation, BINA48 is a transhuman entity, among the first of her kind, an intelligent computer. She is a person embodied in a computer. BINA48, as any human individual, is possessive of consciousness, autonomy, memory, thoughts, beliefs, attitudes, and feelings. However, in lieu of DNA chromosomes, BINA48 was conceived in software. She is currently employed as a tax paying Telecommunications 800 line Customer Service Representative and as an Answers Researcher with Google.



Anthony L. Dutton, Esq. concentrates his practice in the areas of corporate and general business law, including acquisitions and dispositions of businesses. In addition, he has broad experience in representing lenders and borrowers in secured and unsecured lending transactions. He also represents a number of clients with respect to importing and exporting their products and related United States Customs matters.



Martine Rothblatt, J.D. Ph.D. started the satellite vehicle tracking and satellite radio industries and is the Chairman of United Therapeutics, a biotechnology company. She is also the founder of Terasem Movement, Inc.



*Judge David E. Silverman is Judge of the County Court, 18th Judicial Circuit in and for Brevard County, State of Florida.
Elected ~ September, 1998
Term of Office ~ January, 1999 through December, 2010
Assignment: Division IV, General County Jurisdiction (Criminal and Civil)
Brevard County Courthouse, Melbourne
Brevard County Administrative Judge, 2002 through 2004*

Case Citations and links:

Association of Data Processing Service Organization, Inc. v. Camp, 397 U.S. 150 (1970) 87.

[Link to Citation](#)

March 16, 2006 4:02 PM, EST.

Jones v. Alfred Mayer 392 U.S. 409 (1968)

[Link to Citation](#)

March 17, 2006 2:20PM EST

United States v. Guest 383 U.S. 745 (1966)

[Link to Citation](#)

March 17, 2006 12:01 PM EST

¹ Standing: A person's right to bring (start) or join a lawsuit because he or she is directly affected by the issues raised. Oran, Daniel. LAW Dictionary for Nonlawyers 4th Edition. New York: Delmar, 2000.