

Uploading and Branching Identity

Michael A. Cerullo

Received: 11 July 2014 / Accepted: 18 November 2014 / Published online: 2 December 2014
© The Author(s) 2014. This article is published with open access at Springerlink.com

Abstract If a brain is uploaded into a computer, will consciousness continue in digital form or will it end forever when the brain is destroyed? Philosophers have long debated such dilemmas and classify them as questions about personal identity. There are currently three main theories of personal identity: biological, psychological, and closest continuer theories. None of these theories can successfully address the questions posed by the possibility of uploading. I will argue that uploading requires us to adopt a new theory of identity, psychological branching identity. Psychological branching identity states that consciousness will continue as long as there is continuity in psychological structure. What differentiates this from psychological identity is that it allows identity to continue in multiple selves. According to branching identity, continuity of consciousness will continue in both the original brain and the upload after nondestructive uploading. Branching identity can also resolve long standing questions about split-brain syndrome and can provide clear predictions about identity in even the most difficult cases imagined by philosophers.

Keywords Personal identity · Uploading · Consciousness

Introduction

How can we ever really know that we are the same person now that we were yesterday or will be tomorrow? Does the self persist when we wake from a

M. A. Cerullo (✉)

Cincinnati Institute for Cognitive Science, 3898 Drakewood Drive, Cincinnati, OH 45209, USA
e-mail: Cincinnati.psychiatry@gmail.com

dreamless sleep or is our sense of continuity a mere illusion? Could I continue to exist after my brain dies? Such questions about identity have been a major theme throughout the history of western philosophy and are classified as questions about personal identity (Parfit 1984; Olson 2010). Derek Parfit's book, *Reasons and Persons* (Parfit 1984), remains one of the most thorough examinations of personal identity. Parfit explores identity through thought experiments where the brain is split, fused, modified by neurosurgery to change or swap memory and personality, and finally teleported across space to mars. While all of Parfit's examples were theoretically possible, they seemed to be nothing more than science fiction fantasies when the book was written in 1984. The entire debate could be dismissed as esoteric philosophy without relevance to the real world at that time or in the foreseeable future.¹

Now fast forward a mere 30 years to the present and the exponential growth of technology has caught up with science fiction. An understanding of personal identity is now urgently needed to inform critical decisions our civilization must make in the twenty-first century. While teletransporters and radical manipulations of the brain are just as fantastic now as they were in 1984, brain uploading is rapidly becoming a reality. Also called whole brain emulation, uploading involves transferring the informational structure of the brain into a computer (Sandberg and Bostrom 2008; Eth et al. 2013; Sandberg 2013). The European Union has invested over 1 billion Euros over the next 10 years in the Human Brain Project which began in 2013. The project's initial goal is to scan, upload, and emulate a complete mouse brain within 5 years (<https://www.humanbrainproject.eu/strategic-mouse-brain-data>). It will require much more computer power to emulate a complete human brain but there is no difference in principle between uploading a mouse or a human brain. The Human Brain Project aims to use the knowledge gained from the mouse emulation to scan and upload parts of the human brain within 10 years (<https://www.humanbrainproject.eu/strategic-human-brain-data>) and the ultimate goal of the project is to emulate the complete human brain (http://tierra.aslab.upm.es/documents/projects/HBP_flagship_report_for_Europe.pdf). Estimates range from 50 to 100 years before human uploading is possible and depend on how much biological detail is required to emulate the brain (Kurzweil 1999, 2005; Sandberg and Bostrom 2008; Eth et al. 2013). Knowledge of synaptic and dendritic connectivity should allow for a complete emulation of brain function. If necessary, further details at the molecular level (e.g. neurotransmitter and ion concentrations) could be provided to complete the emulation but would require increased computational power (Sandberg and Bostrom 2008; Eth et al. 2013).²

Unfortunately (or perhaps fortunately), we don't have 50–100 years to ponder the philosophical implications of uploading. The technology of whole brain preservation is rapidly advancing and it is estimated that we are only a decade away from being able to permanently preserve the complete human brain (Knott et al. 2008;

¹ Just what is meant by “the future” or “near future” can be ambiguous so it would be more accurate to say very few people believed these issues would ever become practical during their lifetimes.

² There is no evidence that information is necessary at the quantum level (Kandel et al. 2000; <http://www.fhi.ox.ac.uk/brain-emulation-roadmap-report.pdf>; Hayworth 2010).

Hayworth 2012; <http://www.brainpreservation.org/content/overview>). Currently it is possible to chemically preserve a small mammalian brain in enough detail to capture all the details that neuroscience tells us matter for the brain's role as an information processor and which would therefore matter for identity and consciousness (Knott et al. 2008; Hayworth 2012).³ Once a human brain is preserved it could be stored for centuries without any loss of information. When uploading becomes possible in 50–100 years these brains could be taken out of storage and uploaded.

Given these rapid advances in neuroscience it may soon be possible to choose brain preservation at (or right before) death. This procedure should be relatively cheap, about as much as minor surgery and much less than the average cost of a funeral and thus widely available to anyone who wants it (at least in industrial nations initially). We are therefore at a unique point in history where we need to make decisions about the future of humanity based on our best understanding of the philosophy of mind and consciousness. How much money should we invest in developing brain preservation and uploading technology? Individually, should we opt to preserve our brains for later uploading? Will we awake from uploading as we do from a deep sleep or will the awakened individual simply be a simulacra with a unique consciousness (or no consciousness) and fake memories?

In the remainder of the paper I will argue that brain uploading will preserve personal identity. Traditional views of personal identity cannot cope with the possibilities created by advances such as uploading. Instead, we are driven to accept the possibility that personal identity can branch into multiple copies, each maintaining a continuity of consciousness with the original. This is actually the simplest explanation for split-brain syndrome but has been ruled out in the past as being absurd (i.e. unintuitive). This paper will explore the case for branching identity and show how it can resolve the dilemmas that are unsolvable with the three traditional approaches of personal identity.

Theories of Identity

In the context of uploading it is clear by what we mean by personal identity: when I go to sleep to have the uploading procedure, will it be me that wakes up inside the computer? Will uploading be a continuation of consciousness rather than the oblivion of death? In philosophy this type of identity is termed numeric identity and contrasted with qualitative identity. Two things are qualitatively identical if they are indistinguishable copies (e.g. a clone). Numeric identity requires them to be *exactly* the same thing. This is not the best terminology and can lead to much confusion as

³ The best current methods of brain mapping involve scanning thin slices of the brain with an electron microscope. When these slices are created with a diamond knife a resolution of around 50 nm is possible. A newer technique, Focused Ion Beam Scanning Electron Microscopy (FIBSEM), can scan tissue at resolutions approaching 5 nm (Hayworth 2012). This is enough detail to capture a complete map of all synaptic and dendritic connections and their strengths but not the details at the molecular level (neurotransmitter and receptor levels) (Hayworth 2012). However, current brain preservation techniques preserve information at the molecular level and there are several potential techniques in development that can provide molecular level scanning resolution (Sandberg and Bostrom 2008).

we will see later in this paper. For our purposes it is better to define personal identity as the kind of identity that preserves the continuity of consciousness and that is how we will use the term throughout this paper. We can then define continuity of consciousness as the continuation of subjective experience within the same entity through time.

There is a large body of literature on personal identity with its own terminology and little consensus (Shoemaker 1984; Swinburne 1984; Parfit 1984; Olson 2010). In a recent paper Chalmers reviewed the major philosophical stances on personal identity and we will adopt his terminology (Chalmers 2010). There are three main theories of personal identity: biological, psychological, and closest continuer. Those who hold the biological theory believe that the continuity of physical brain is essential for identity and must continue for there to be continuity of consciousness. In contrast, the psychological theory requires psychological continuity for identity to be preserved: it is the information content of the brain that matters and as long as memory and causal structure are recreated then identity should continue. Those who hold the closest continuer theory claim that consciousness will continue in whatever entity is most identical to the original. Each of these three theories has many variants, most of which share the essential points just discussed.

The Failure of the Standard Approaches to Identity

The standard theories of personal identity cannot provide satisfactory answers to questions about the continuity of consciousness in uploading. Yet we do not need to discuss future technologies to see the failure of the standard theories; they are unable to deal satisfactorily with split-brain syndrome. After a corpus callosotomy for intractable epilepsy, patients appear to have a splitting of consciousness into two selves, each of which is contained in a different brain hemisphere. This has been termed split-brain syndrome (Gazzaniga et al. 1962; Gazzaniga 1967).⁴ The implications of split-brain syndrome have been discussed since the syndrome was first discovered yet there is no consensus on what happens to personal identity in each brain hemisphere (Gazzaniga 1967; Nagel 1971; Parfit 1984; Shoemaker 1984; Swinburne 1984). We will adopt Parfit's terminology and call the consciousness created in the left hemisphere Lefty and the consciousness created in the right hemisphere Righty (Parfit 1984).

Let's examine how each of the three theories of personal identity deal with split-brain syndrome. Before we can do this we need to examine individuals with a single brain hemisphere. People born with a single brain hemisphere do not suffer any loss of consciousness. Others lose one hemisphere through disease or accident and it appears their continuity of consciousness continues (Pulsifer et al. 2004). In fact,

⁴ It does not appear that separation into hemispheres is required for high level cognition or consciousness (Levy 1977; Vallortigara and Rogers 2005; MacNeilage et al. 2009). Rather it appears to be a product of the particular evolution of the early ancestors of vertebrate animals. If human brains hadn't evolved in this particular manner then split-brain syndrome would likely not be possible. The octopus, which diverged from vertebrate animals for 750 million years ago, has a completely different brain architecture without brain hemispheres.

neurosurgeons will remove one hemisphere to save someone's life if no other options are available in a procedure termed a lobectomy (Oncel et al. 2007; Pulsifer et al. 2004; Hayworth 2010). Given the empirical facts it would be extremely difficult for any theory of personal identity to deny continuity of consciousness after a lobectomy. Cases of lobectomy also provide information on how much brain, psychology, or closeness is sufficient for continuity: half or more.⁵ Biological theories would claim that the hemisphere is a survivor of the original biological brain and this is enough to keep continuity of consciousness. Psychological theories would claim the surviving hemisphere keeps enough of the person's psychological structure intact for continuity of consciousness. Those who hold the closest continuer theory would state the survivor of a lobectomy has continuity of consciousness as clearly the surviving hemisphere is more identical than anything else existing to the original brain.

A corpus callosotomy severs most of the connections between the two brain hemispheres and is used to treat intractable epilepsy. Neuropsychological testing after the procedure shows that there are now two independent conscious selves (Gazzaniga et al. 1962; Gazzaniga 1967). Which of these two selves maintains continuity of consciousness with the whole brain before surgery? If I have this operation, do I wake up as Righty or Lefty? Remember that if a lobectomy is performed then the sole surviving hemisphere, whether Righty or Lefty, maintains continuity of consciousness. The best the biological theory can do to cope with split-brain syndrome is to claim that whatever hemisphere contains the most biological similarity to the original brain will maintain continuity of consciousness. Yet each brain hemisphere is equally biologically continuous with the whole brain. Similarly, adherents of the psychological theory could claim that whichever hemisphere is most psychologically continuous with the original brain will maintain continuity of consciousness. Again, both hemispheres appear equally psychologically continuous with the original brain and each maintains continuity of consciousness after a lobectomy.⁶ Therefore the biological and psychological theories cannot answer the question of whether I wake up Righty or Lefty.

Supporters of biological and psychological identity theories have proposed two possible solutions to the problems posed by split-brain syndrome. First, they can claim that neither Righty nor Lefty retains continuity of consciousness. There several problems with this solution. Each hemisphere retains continuity of consciousness when it is the only surviving hemisphere, so why would this change if the other hemisphere is not destroyed? The continuation of consciousness in one hemisphere should not depend on what happens to an independent object. To complicate this problem even further, we could separate the hemispheres and keep one hemispheres,

⁵ Lobectomy cannot tell us what the minimum amount of brain, psychology, or closeness is necessary for continuity of consciousness.

⁶ Because brain function is asymmetric and since the left brain most often controls language, it could be argued that Lefty is the more psychological half. Yet this bias towards language seems unjustified as it appears that both hemispheres contribute equally (if differently) to the original self. In addition, in some individuals brain function is symmetric and language function is equally divided between hemispheres (Galaburda et al. 1990; Springer et al. 1999; Corballis 2009). Therefore brain asymmetry cannot be used to rescue psychological (or as we will see closer continuer) identity.

say Lefty, in cryogenic suspension. Would Righty retain continuity of consciousness because it is the only hemisphere currently conscious: or would the mere potential of reactivating Lefty disrupt the continuity of consciousness of Righty? If Righty does retain continuity of consciousness, what happens if we then destroy Righty and reactivate Lefty? Does Lefty then have continuity of consciousness? Biological and psychological identity cannot provide any non-arbitrary predictions in these cases. Therefore claiming neither hemisphere maintains continuity is not a satisfactory solution. The second possible solution is to assume that the brain was originally composed of two separate selves, each with their own continuity of consciousness, and these two selves are separated by the corpus callosotomy. This solution fails because it appears ad hoc, and more importantly, we experience only one unified conscious self (and it is this unified self that we set out to explain in the first place). The possibility that there was only one self to begin with and that both hemispheres retain continuity of consciousness with this self is not considered or is ruled out arbitrarily by non-branching axioms (Swinburne 1984; Olson 2010).

The closest continuer theory is the only identity theory that clearly sets up criteria to determine the continuity of consciousness after fission: only the part of the brain most identical to the original brain maintains continuity of consciousness. Yet if the brain is split into two equal parts the closest continuer theory is not able predict which half maintains identity. This appears to be the case in split-brain syndrome and the closest continuer theory can only arbitrarily decide which half (or if neither) maintains continuity of consciousness. Predictions made with the closest continuer theory will also depend on what we do in the present or future with the other hemisphere.

Parfit rejects the standard theories of identity and develops his own unique solution to the split-brain dilemma: what matters in personal identity is the continuation of what he terms the relationship R (Parfit 1984). Parfit defines R as a continuation of memory and psychology. What differentiates R from psychological identity is that Parfit rejects any notion of non-reductive identity. Parfit claims that memory and psychology are entirely what a person consists of and he rejects anything extra such as consciousness or qualia. Parfit's theory seems to deal successfully with cases of lobectomy or corpus callosotomy: the surviving hemisphere in the case of a lobectomy, or Righty and Lefty in a corpus callosotomy, all continue a significant portion of the person's original psychology and memory and are therefore forms of survival. The problem with Parfit's solution is that he rejects the very notion of continuity of consciousness. His reductive physicalism denies the existence of qualia and without qualia there is no continuity of consciousness. Remember, our goal is to find out if our stream of consciousness continues when we wake up after the uploading procedure. Parfit is saying we were never awake to begin with. To understand personal identity we cannot accept reductive physicalism. I do not believe we can solve dilemmas about identity without accepting the reality of the hard problem of consciousness and accepting qualia (Chalmers 1995a, 1996).⁷

⁷ We can accept the hard problem and still agree that everything mental supervenes on the physical. We just deny that qualia are identical with matter and instead assert that they are a basic but non-physical part of the universe.

Machine Consciousness

We need to take a brief diversion at this point and talk about machine consciousness. There would be little point to uploading if it was not possible for digital computers to be conscious. The prevailing viewpoint in contemporary western philosophy of mind after behaviorism (which was essentially a form of eliminativism about consciousness) was the identity theory. The identity theory states that the mind (and consciousness) is identical to brain states. Hilary Putnam's argument of multiple realizability dethroned the identity theory and functionalism took over (Putnam 1967, 1988).⁸ Functionalism asserts that it is the function and causal structure of the brain that is responsible for consciousness and it remains the dominant theory of mind today (Chalmers 1995a, 1996; Levin 2013). It is the cognitive structure of the brain (i.e. information) rather than the physical matter that counts; therefore any machine that duplicates the cognitive architecture of the brain will be as conscious as the original brain.

Chalmers introduced the fading/dancing qualia argument to support functionalism and we will review it in detail here as it will reoccur often in subsequent arguments (Chalmers 1995b, 1996). Qualia (or the plural form quale) refer to specific conscious experiences such as perceiving the color red. More complex phenomenal experiences also have an associated qualia and there is a specific qualia for the overall unified state of consciousness a person is experiencing at any one moment. To understand the fading qualia argument, imagine replacing the neurons in a person's brain one by one with nanomachines that are able to perform all the functions of the neuron. Then remove all elements of the nanomachines except their interface with other neurons. The elements of the nanomachines removed could be emulated in a computer which could radio instructions to the machine interfaces connected to the remaining biological brain. As can be seen, this process will slowly transfer the function of the brain into a computer. Since the brain is destroyed during the process we call this destructive uploading. The process is also happening over a period of time which is termed gradual uploading. The cognitive architecture and thus functional structure of the brain never change as the brain is slowly replaced. Therefore if functionalism is true consciousness remains unchanged throughout the process. If functionalism is false, then as the brain is slowly replaced consciousness would begin to fade away. Yet by the design of our experiment the function of the brain never changes. The person's consciousness would slowly fade away and yet this could not affect their behavior in any way (e.g. they couldn't say "I am starting to feel a little odd...").⁹

⁸ Multiple realizability asserts that it is possible for the same mental state to be realized by different physical states.

⁹ Fading and dancing qualia violate the principle of structural coherence (Chalmers 1995b, 1996). Structural coherence states that:

Whenever there is a conscious experience, there is some corresponding information in the cognitive system that is available in the control of behavior, and available for report and global control. Conversely, it seems that whenever information is available for report and global control, there is a corresponding conscious experience.

(Chalmers 1995b).

A slight variation of our experiment can create the dancing qualia argument (Chalmers 1995b, 1996). Instead of destroying the neurons, the nanomachines keep them intact and reversibly take over their function. Once the nanomachines are in place they can be switched off to allow the original neurons to take over and then switched back on whenever we choose. A significant part of the brain can be duplicated (e.g. all regions responsible for vision) and the machines can then be turned on and off repeatedly. According to functionalism nothing will change in consciousness as the function of the brain is unchanged. If functionalism is incorrect, a significant aspect of consciousness will disappear when we turn the machines on. A moment later when we turn off the machines and let the neurons resume their function, the missing aspect of consciousness will immediately reappear. Therefore the person will experience these qualia turning on and off. Yet these dancing qualia could not affect the person's behavior in any way (e.g. they can't say "wow, I just went blind for a minute but then luckily my vision came back").¹⁰ The same process that created the dancing qualia effect could also be used to create a digital duplicate of the entire brain while keeping the physical brain intact. This is termed nondestructive uploading.

There are many additional variations on the fading/dancing qualia argument, all of which illustrate problems with the identity theory and provide support for functionalism. One variation is lacuna qualia where we once again replace a part of the brain with our nanomachines. Continuing with our previous example, let's replicate the brain regions responsible for vision. Now, instead of turning that region on and off, we turn the nanomachines on permanently and destroy the original neurons. The person would continue to perform perfectly on all visual tests and would claim to have normal vision. Yet the identity theory would assert that the person lost all visual qualia the moment the nanomachines were turned on. Just as in fading or dancing qualia, the claims made by the identity theory seem extremely implausible. We can create floating qualia by replacing most of visual regions of the brain with nanomachines but keeping the regions responsible for color vision. Again the person continues to perform perfectly on all visual tests. The identity theory would assert they only experience color qualia and no other vision. Once again our subject cannot act on or report these colors that float free of any other visual experience. Similar examples could be created using any other type of qualia. For example, we could replace the neurons responsible for a specific memories, mathematical abilities, political beliefs, and so on. We could thereby interrupt the unity of conscious experience at any point by removing any specific qualia we desire. These qualia thought experiments can be generated indefinitely and all show the absurdities that would need to be accepted if the identity theory is true. They present no obstacles to functionalism.

Uploading and the Standard Theories of Personal Identity

If functionalism is correct, then digital brain emulations will experience the same consciousness as biological brains. Yet functionalism does not address questions of

¹⁰ The phenomenal experience of temporary blindness is a common occurrence in Transient Ischemic Attack and has its own set of qualia and corresponding neural correlates of consciousness.

personal identity and cannot help us determine if an upload maintains continuity of consciousness with the original brain (Chalmers 2010). If I go to sleep at my local upload center in the year 2080, will I wake up in a computer simulation or will a newly created conscious entity awake for the first time complete with false memories? What will happen if I choose nondestructive uploading and keep my original self intact? The three traditional views of personal identity cannot help us answer these questions. Many of the same problems encountered in split-brain syndrome recur and there are also new dilemmas unique to uploading.

Psychological identity asserts that a destructively uploaded brain retains continuity of consciousness since it maintains psychological continuity. If we nondestructively upload and keep the original brain intact, psychological identity cannot tell us which of the two retains continuity of consciousness (recall the psychological theory has the same problem when both Righty and Lefty survive). The closest continuer theory would also state that a destructively uploaded brain retains continuity but is similarly unable to determine whether the upload or biological brain retains continuity of consciousness in nondestructive uploading.¹¹ The psychological and closest continuer theories can also claim that neither the brain nor upload maintain continuity of consciousness after nondestructive uploading. However, it seems ad hoc to claim that where there are two equally identical psychological or closest continuers of a brain, neither retains continuity. This could be added on as an extra axiom but it does not logically follow from either theory. Claiming neither the brain nor upload maintain identity leads to the same problems the theories encountered when denying continuity to Right and Lefty: once again continuity of consciousness in one entity depends on what happens in space and time to a completely independent entity. The last option would be to claim that both the brain and the upload maintain continuity of consciousness. This is the only solution that can make specific predictions about identity that avoid dependence on what happens in space and time to an independent entity. Yet this possibility is often ruled out as absurd and many philosophers add a non-branching axiom stating identity can never branch (Swinburne 1984; Olson 2010).

We have just shown that psychological and closest continuer theories are unable to make clear predictions in nondestructive uploading as long as they adopt the non-branching axiom. It can also be shown that that the two theories lead to absurd conclusions if they accept the non-branching axiom. There is a strong affinity between these theories and the identity theory of mind we reviewed previously when we defended functionalism. Psychological and closest continuer identity both claim that continuity of consciousness is *identical* to a particular property.

This identity claim leaves both theories open to a type of fading/dancing qualia argument that we will call popping qualia. We will examine popping qualia arguments in each of the four possible outcomes of nondestructive uploading: only

¹¹ Two options are open to a proponent of the closest continuer theory. First, they could claim the upload and brain are equally close to the original, in which case the theory breaks down as it allows only one successor by definition. On the other hand, our proponent could claim the brain maintains continuity as it doesn't change while the upload is a newly created copy. In this case the closer continuer theory reduces to biological identity and inherits the problems with that theory.

the brain maintains identity; only the upload maintains identity; neither brain nor upload maintains identity; and both brain and upload maintain identity.

To invoke the popping qualia argument, first gradually nondestructively replace all the neurons in the brain with nanomachines. Now assume only the biological brain maintains continuity of consciousness. Suppose we turn off the neurons as they are replaced by the nanomachines. At some point in this process when greater than half the brain has been replaced (from split-brain syndrome) the continuity of consciousness will suddenly terminate. If functionalism is correct, a new conscious entity will be created at this point in the brain/machine hybrid with false memories. This new entity will be functionally identical to the original and have the exact same conscious experiences. Once the brain is completely replaced by our nanomachines, we start turning the neurons back on one by one. The newly created consciousness could continue during this process or be replaced by another conscious entity. Neither the psychological or closest continuer theory makes clear predictions in this building up of a brain scenario. For the current argument it doesn't matter if a third conscious entity is created or not. To invoke the fading/dancing qualia argument, we repeat this brain replacement and reinstatement at an accelerated pace. There is no reason in principle why we could not complete one step (replacement or reinstatement) every 30 ms (or less). This time period is less than the minimum time increment of conscious perception (Efron 1970; Ruhnau 1995; Breitmeyer and Ogden 2000; Pockett 2003; Coren et al. 2004; Dainton 2014). Psychophysical studies also show that events that occur within 40 ms or less of one another are perceived as simultaneous (the exact time varies by stimuli type) (Efron 1970; Ruhnau 1995; Breitmeyer and Ogden 2000; Pockett 2003; Dainton 2014).

Returning to our example, two things could be happening as we repeatedly replace and reinstate the brain. The first possibility is that a new conscious entity could be appearing within the minimum phenomenal time span of consciousness. This would create a bizarre form of popping qualia where the newly created entity bursts into existence in the middle of the minimal temporal phenomenal experience. In this case either the experience is shared between two conscious entities (the last moment of one and the first moment of the second) or the new conscious entity begins at the next moment of experience. Neither option seems desirable. Yet things get much worse when we consider the effect popping qualia have on the stream of consciousness. Because our stream of consciousness extends for several seconds in working memory, the new conscious entity would understand that the last several seconds of time was not experienced by them. By the design of our experiment the functional structure of the brain/machine hybrid is identical at all times and would not deviate in any way from the original brain. Therefore the subject could not act on these bizarre popping qualia experiences and we have the same absurd situation created in the fading/dancing qualia arguments. In addition, many complicated concepts take the full capacity of working memory to understand. If a new conscious entity was created in the middle of understanding a difficult sentence or mathematical equation, they would seem to have access to only half the needed awareness. Once again there would be no behaviorally detectable difference and these missing qualia would be another violation of Chalmers' principle of structural coherence.

The second possibility of our replacement and reinstatement of the brain example is that we are already Nagel millisecond-people who only exist for one quantum of phenomenal experience and therefore nothing would change in our experiment.¹² In this case personal identity has always been an illusion. So we are left with either popping (and missing) qualia or an extreme skepticism about identity, neither of which seems plausible. The popping qualia argument would be identical if we were to claim only the upload maintains identity in nondestructive uploading. It can be easily seen that the argument also works if we were to claim neither brain nor upload maintains identity. The argument *does not* work if we claim both the brain and upload maintains continuity of consciousness. Therefore we have shown psychological and closest continuer theories lead to absurd conclusions if they include non-branching axioms.

We have been focusing on the psychological and closest continuer theories but we can apply many of the same arguments to the biological theory of identity. Biological identity is the only theory that seems to give a determinate answer in nondestructive uploading: only the original brain will retain continuity of consciousness. Yet biological identity shares the same links with the identity theory as the psychological and closest continuer theories and is therefore also vulnerable to popping qualia arguments.¹³ Proponents of biological identity would claim that only the brain maintains identity in nondestructive uploading and we just saw that this leads to popping and missing qualia. Biological identity would also deny continuity of consciousness to a destructive upload and it can easily be seen that this would also lead to popping qualia.

Branching Identity and Intuition

At this point we will diverge from the three standard theories of identity. As we indicated previously, philosophers use the term numeric identity when talking about the continuity of consciousness. The problem with the term numeric identity is that it assumes by definition that there can only be *one* numerically identical thing. Once again we need to keep our final goal in mind: does consciousness continue when we awake from uploading? We should avoid biasing ourselves from the outset against an entire class of solutions to the identity problem. Branching identity theories are free from this bias and allow continuity of consciousness to split (or branch) and continue on in more than one self. Branching identity is often dismissed as absurd or impossible and we have seen that many theories include arbitrary non-branching axioms (Swinburne 1984; Olson 2010). In the next section we will examine how branching identity can resolve many of the dilemmas created by split-brain

¹² Nagel coined the term day-person to refer to the possibility that each evening after we go to sleep we die and a new conscious being with false memories awakes the next day (Parfit 1984). There is of course no way to rule out that we are all day-people, or even that we are minute-people that are born anew each minute with false memories.

¹³ Biological theories of personal identity are also vulnerable to the fact that the matter in our brain is being constantly recycled. This challenges the idea that our brains maintain a continuity of matter throughout our lives that is tied to personal identity.

syndrome and uploading. But before examining the benefits of branching identity we need to tackle the arguments that it is impossible or absurd.

I concede that branching identity is unintuitive. In our evolutionary history problems involving personal identity were not relevant to survival. Thus we have no built in cognitive mechanisms to model personal identity. In a similar way, while our brains do have mechanisms to model three dimensional Euclidean space we have no such built in ability to perceive higher dimensions. Without these innate abilities we have no intuition about higher dimensions. Yet modern physics tells us that we live in four dimensional space–time. We have been able to generalize our understanding of geometry and abstractly develop and understand models of higher dimensions. The first lesson here is that we should use caution in ruling something out because it seems unintuitive to our primate brain. The second lesson is that it may possible to extrapolate from our experiences of personal identity and gain some intuition about branching identity. In fact, it is probably much easier to imagine branching identity than the fourth dimension.

Thinking about alternate history (or what philosophers call counterfactuals) is thinking about branching identity. We seem to have no trouble understanding discussions of how history would change if certain things were altered. There are survival advantages to being able to think about alternatives to our prior behavior and thus we shouldn't be surprised at our ability to do so. When we think about these possibilities we are imagining another version or branch of ourselves. That version branches off from us at the point where our behavior diverged. We seem to be able to easily grasp that the alternate self and our current self shared continuity of consciousness the moment before they began to diverge. This is branching identity, and it clearly isn't as unintuitive as we may have initially supposed. Of course one might object that the real challenge is to imagine other versions of ourselves existing at the same time rather than in alternate realities. Once again we do this all the time: there is a vast collection of fiction discussing time travel. We have no problem following a story where someone travels back in time and meets an earlier version of themselves. In these stories we easily imagine two people who share a continuity of consciousness existing at the same time.

These examples can help us start to imagine branching identity. It is probably beyond our ability to completely envision continuity of consciousness continuing in both Lefty and Righty or the brain and the upload. After all, this seems to be like imagining that two things are one. But consciousness and qualia may not behave like familiar physical objects and thus our intuition about objects may mislead us when thinking about the continuation of subjective identity. Again returning to the fourth dimension, we can use our knowledge of Euclidean geometry to understand the fourth dimension from an abstract perspective but will never be able to truly see a hypercube in our imagination. In a similar way we will probably never be able to completely imagine consciousness splitting off into parts that are equally ourselves. Luckily our goal wasn't to show branching identity is obvious: instead we just wanted to demonstrate that we have the ability to abstractly understand it and thus counter the arguments that it is absurd or impossible. Now are free to examine the arguments in support of branching identity.

A Definition of Branching Identity

Earlier we stated that branching identity theories allow the continuity of consciousness to branch and continue in more than one self. This is the beginning of a definition but there are still many details to specify. First, we need to examine how much continuity in psychology or biology is needed to make a branch. If we upload a small piece of the temporal lobe we don't want our definition to assert that this small upload has continuity of consciousness. While possible, it seems highly unlikely that such small parts of the brain have continuity of consciousness with the whole brain. We just don't know enough about consciousness to commit ourselves to these kinds of details at this point. The exact amount of psychological continuity needed for identity is an empirical question that may not be answerable until we have a more fundamental theory of consciousness. However, using what we do know about split-brain syndrome, a conservative definition would be that continuity of consciousness branches when at least half a mind (biological or uploaded) is split off from the original. We also need to specify what kind of continuity matters: biological or psychological. There could be a biological and a psychological version of branching identity (the closer continuer theory rules out branching by definition). Biological branching identity would claim that only the fission (or cloning?) of half or more of a biological brain would create branching identities. Psychology branching identity would allow for branching in any entity that continued half or more of the psychology of the original entity (brain or upload). We can rule out branching biological identity because it would deny continuity of consciousness to uploads and would therefore be vulnerable to popping qualia arguments. Thus psychological branching identity is the only remaining option and we provide a formal definition as follows:

There will be continuity of consciousness between any conscious entity P_1 at time T_1 and P_2 at time T_2 if P_2 contains half or more of the psychological structure of P_1 at time T_1 and is activated at any time T_2 subsequent to T_1 .

This definition leaves open some difficult questions concerning the relationship between branching identity and time; e.g. will 2 uploads created 1 min apart share personal identity? These are difficult questions best avoided until we have a better fundamental theory of consciousness. We will therefore avoid making any commitment to these questions by setting $T_1 = T_2$. The simpler definition of instantaneous psychological branching identity becomes:

There will be continuity of consciousness between any conscious entity P_1 at time T_1 and P_2 at time T_1 if P_2 contains half or more of the psychological structure of P_1 at time T_1 and is activated at time T_1 .

One unstated assumption of both definitions is that the new entity P_2 is not made up of other cognitive structures. Uploading would certainly allow merging of multiple conscious entities but again this type of complication is best avoided until we have a more fundamental theory of consciousness. Hence our definitions assume that P_2 is only made from components of P_1 .

The Benefits of Branching Identity

Instantaneous psychological branching identity allows detailed predictions in split-brain syndrome and uploading (for the sake of brevity we will refer to this theory as branching identity). Branching identity predicts the empirical result seen in split-brain syndrome: both Righty and Lefty appear to have continuity of consciousness with the whole brain. It also predicts that continuity of consciousness will be preserved in nondestructive instantaneous uploading. In this case two identical entities would emerge, the brain and the upload, each sharing continuity of consciousness with the original. In cases of destructive uploading, there would be only one branch and branching identity would reduce to psychological identity and would predict that destructive uploading preserves continuity of consciousness.

Branching identity can also give specific predictions in the many complicated scenarios created by machine consciousness. For example, once a mind is uploaded it is a trivial task to make as many copies as desired. Each of these copies would share continuity of consciousness with the original upload at the time it was copied. Finally, branching identity gives a clear prediction in one of the most bizarre scenarios discussed in personal identity, complete memory transfer (Swinburne 1984; Unger 1990). In this scenario, all of person A's memories and psychology are transferred into person B's brain. This transfer can be done destructively or non-destructively. We can call the new person with B's biological brain and A's memories (and psychology) B_A . It is not entirely clear if this example is coherent, but if it is branching identity predicts that there is continuity of consciousness from A to B_A in the case of destructive memory transfer and from A to both A and B_A in nondestructive memory transfer.

As can be seen, branching identity gives specific predictions in even the most challenging cases imagined by philosophers. No other theory of personal identity can match this. These predictions do not depend on anything besides the brains (or uploads) in question. Thus branching identity seems the most parsimonious solution to the difficult challenges posed by philosophers. It also avoids the many problems including popping qualia that occur with the three traditional theories of personal identity. The main obstacle to branching identity is accepting that consciousness can branch. Yet we already have very strong empirical evidence that this occurs in split-brain syndrome. We have also seen that branching identity, while not part of our everyday experience, is not absurd or impossible. We will make one final direct argument in the next section supporting branching identity and this will involve the concept of qualia space.

Qualia Space

Qualia space refers to a mathematical space that encompasses all qualia (Stanley 1999).¹⁴ For every possible qualia there is a unique point in this space. Qualia space

¹⁴ The geometry of qualia space is "a closed pointed cone in an infinite-dimensional separable real topological vector space" (Stanley 1999).

includes not only individual qualia such as the sensation red but also more complex qualia created by combining individual qualia. Since our consciousness experience is a unified whole it maps onto a unique single point in qualia space.

There are two main interpretations of qualia space. The first is that qualia space has a real existence outside space–time. This seems to parallel the belief many mathematicians have that mathematical structures exist in a reality outside space–time. Since this mathematical view is termed mathematical realism we can call this first position qualia realism (technically this could be quale realism but qualia realism sounds better). The other possible interpretation of qualia space is that it is just a potential space that maps all the possible quale that could exist. Since this parallels mathematical constructivism we can term this position qualia constructivism. Luckily we don't need to decide between these possibilities as either interpretation still supports our argument for branching identity.

To see how qualia space supports branching identity, we need to consider two facts. First, by definition every possible phenomenal experience has an associated qualia. Therefore the experience of having continuity of consciousness through time also has an associated qualia. The second fact to consider is that every possible qualia can be mapped to a unique point in qualia space. Putting these two facts together it can be seen that the experience of the continuity of consciousness is just another qualia in a person's unified conscience experience. This unified conscious experience is then mapped onto a unique point in qualia space. Once we understand the psychophysical laws linking physics to qualia (the neural correlates of consciousness) we can create an entity that has a conscious experience that maps to any point we desire in qualia space. It will therefore be possible to recreate the *exact* subjective experience of a person at any time point (assuming of course we have a complete map of the information state of the brain or upload at that time). This subjective experience also contains the qualia associated with continuity of consciousness. Any two entities that mapped onto the same point in qualia space would be completely identical and share the same phenomenal experience of personal identity. Therefore any uploading procedure that mapped the upload to the same point in qualia space as the biological brain would have the same continuity of consciousness.¹⁵ Qualia space thus leads us to the conclusion that quale are all there is to personal identity and the continuity of consciousness. Kenneth Hayworth makes a similar proposal in defense of his own version of branching identity (Hayworth 2010). Yet many philosophers would still claim this doesn't capture all there is to identity, and this is our next topic (Parfit 1982, 1984; Swinburne 1984; Chalmers 2010).

Further Facts

Many philosophers claim the three traditional theories of identity (and presumably branching identity) are missing something essential about personal identity. This is

¹⁵ It is better to defer difficult questions about how close points need to be in qualia space to share identity and assume this is true only if they map to the exact same point.

often referred as the “further fact” problem (Parfit 1982, 1984; Swinburne 1984; Chalmers 2010). Swinburne believes that what is left out in the three traditional theories of personal identity is something akin to the soul (Swinburne 1984). Swinburne proposes a version of substance dualism where the soul continues in a person through time and is responsible for personal identity.¹⁶ There are several significant objections to this type of further fact argument. First, these arguments inherit all the limitations of property dualism (Chalmers 1996; Robinson 2012). It is also not clear why the soul would not get transferred by fission. The neurological facts tell us the mind supervenes on the brain, so why wouldn’t the soul move along with the mind? Thus Swinburne’s invocation of the soul does not help him resolve split-brain syndrome. Does the soul go to just Righty or just Lefty? Would the soul disappear altogether after the operation, or would the soul be duplicated so that both Righty and Lefty now have the same soul? Swinburne doesn’t provide any satisfactory answers to these questions.

The claim that something is missing in identity could also be interpreted as a form of the knowledge argument. The most famous knowledge argument was formulated by Frank Jackson to cast doubt on reductive explanations of consciousness and involves Mary the neuroscientist (Jackson 1982, 1986). Mary was a neuroscientist who knew everything there was to know about color vision but had spent her entire life in a black and white room. The knowledge argument states that Mary really doesn’t know everything there is to know about color vision because she has never actually had the conscious experience (qualia) of seeing color. There a large literature on Frank Jackson’s argument and no consensus on whether it is successful or not (Ludlow et al. 2004; Nida-Rümelin 2010).

How would the knowledge argument work in personal identity? Our discussions of identity have assumed the existence of qualia so knowledge arguments in personal identity must be getting at something very different. Returning to Jackson’s example, suppose we have another neuroscientist, Mary’s brother Peter, who knows everything there is to know about personal identity and continuity of consciousness but doesn’t experience it. First, we need to determine if this is a coherent argument. Remember, we are assuming qualia so Peter cannot be a zombie without any phenomenal experience. Yet by design Peter would lack any phenomenal experience associated with continuity of consciousness. Peter would therefore be much worse off than someone with severe memory problems. The most severe memory impairment is the loss of the ability to form short or long term memory and these patients only have access to working memory (Spiers et al. 2001). Any memory deficit more severe than this would be classified as a disorder of consciousness and the person would have severe functional impairment. Peter’s lack of continuity of consciousness would severely limit his short term memory and he would have similar behavior to someone with a severe disorder of consciousness. He would lack the capacity to coherently discuss identity and hence could not know everything there is to know about it and the example falls apart. Popping qualia and

¹⁶ The non-reductive view of qualia I have assumed is a form of property dualism (Chalmers 1995b, 1996). Property dualism assumes there is only one type of substance, matter, and that matter has both physical and mental properties. Substance dualism claims that matter and mind are two different substances (Chalmers 1996).

lacuna qualia arguments also cast doubt on the coherency of the argument. Therefore it seems that knowledge arguments cannot be used to argue for further facts about identity.

Finally, it is possible to connect the further fact argument with theories of biological (or physical) identity. It has been argued (Corabi and Schneider 2012) that if nothing physical is transferred from the brain to the upload, then how could anything (such as identity) move between them? If nothing in the physical world can jump between locations in space without first moving between them, how can personal identity do this? We can name this the No Free Movement Further Fact (NFMFF) argument. If this argument is just another way of stating the biological theory of personal identity, then we have already addressed it earlier in the present paper. If it is another version of the soul argument it was also addressed previously. If it is neither of these then functionalism seems to provide a counterargument. Functionalism states that it is not anything biological or physical that moves between brain and upload but instead information. Information can only be moved via matter so there is no violation of physics in functionalism. If proponents of NFMFF deny functionalism, they are left with the dilemma of fading/dancing qualia. NFMFF arguments may also be based on the perceived difference between fission and copying. Fission seems to preserve some physical continuity while uploading does not. Therefore NFMFF arguments could grant branching identity in split-brain syndrome but not in uploading. This is problematic because the fission/copying distinction seems to be an accident of our biological design. It is easy to see that this distinction disappears in a digitally evolved alien consciousness or an uploaded mind where all that matters is bits of stored information. For these beings fission and copying are identical processes.

Branching Identity and Uploading

If we accept functionalism and the best current neuroscience (e.g. no special quantum process or non-computable processes necessary for consciousness), then branching identity tells us that uploading is an acceptable way to allow for the continuation of consciousness. If the goal of uploading is life extension, then branching identity would suggest we pursue destructive uploading. While this may seem paradoxical at first, i.e. the best way to save the brain is to destroy it, it can be easily explained. After nondestructive uploading at time T_1 , we have two entities, the original brain (P_1) and the upload (P_2). At the next moment in time T_2 when the uploading is finished, P_1 and P_2 will be completely identical and they will both share continuity of consciousness with P_1 at time T_1 . According to branching identity there is no reason to favor P_1 or P_2 as being the “real” you. Qualia space also suggests there is no reason to favor either being P_1 or P_2 as they both map onto the same identical point in qualia space. At time T_3 and beyond however, both entities diverge and go on to have separate conscious experiences and emerge as different selves. If P_1 is allowed to continue after uploading then nothing will have changed for P_1 ; it is still a biological organism and will experience death as before. Therefore destructive uploading is preferable as it avoids this terminal branch.

Chalmers has argued that gradual destructive uploading has the greatest likelihood to preserve continuity of consciousness (Chalmers 2010). The benefit of gradual destructive uploading is that the person could remain conscious during the entire process and thus presumably also maintain continuity of consciousness throughout. Since no branching occurs in gradual destructive uploading, branching identity reduces to psychological identity and would also predict identity is preserved. The problem with gradual destructive uploading is that it depends on advanced brain replacing nanotechnology. While this technology is plausible in principle it will most likely only be developed long after destructive uploading via brain preservation and serial scanning has been perfected. If we accept branching identity, there is no difference between gradual and instantaneous destructive uploading. Kenneth Hayworth reaches the same conclusion and thus advocates spending our resources on the more promising techniques that could allow destructive uploading in the near future (Hayworth 2010).

Conclusion

If we fully accept the empirical observations in split-brain syndrome we are led away from traditional views of identity. Instead, we are drawn towards a theory of identity where the continuity of consciousness can continue in multiple branches. Contrary to the assumptions of past philosophers, there is nothing incoherent or absurd about branching identity. We have limited intuition about things far removed from our day to day experience. Common sense physics has been overthrown by quantum mechanics and relatively. In a similar way we need to expand our views of personal identity. When we closely examine the possibility of branching identity it is not as unintuitive as we might initially suppose and can be abstractly understood in the same way we can come to understand modern physics.

When discussing consciousness we can never expect complete certainty. We cannot know with certainty that solipsism is false or that we are not day-people who die each night when we fall asleep. Yet our common experience has led us away from such views. Questions about personal identity are empirical questions, and just like any other fact about the universe and they will always be open to revision. Yet our best current understanding leads us to branching identity. The empirical evidence (split-brain syndrome) supports branching identity. Only branching identity can provide specific predictions in even the most outlandish scenarios. Finally, popping and fading/dancing qualia arguments support branching identity.

Uploading has the potential to change the way we understand ourselves and our place in the universe. The breathtaking pace of technology has brought us to the point today where all the technology necessary for uploading is now feasible. If the progression of technology continues and animal experiments demonstrate the feasibility of uploading then this should be viewed as life extension technology. Brain preservation and later destructive uploading will preserve continuity of consciousness. The rational choice is to spend whatever resources are necessary to understand, develop, and apply this technology to those who choose to use it. The answer to our ultimate question is yes, it will be you that wakes up inside the computer.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Breitmeyer, B., & Ogden, H. (2000). Recent models and findings in visual backward masking: A comparison, review and update. *Perception and Psychophysics*, *62*, 1572–1595.
- Chalmers, D. (1995a). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, *2*(3), 200–219.
- Chalmers, D. (1995b). Absent qualia, fading qualia, dancing qualia. In T. Metzinger (Ed.), *Conscious experience* (pp. 309–328). Imprint Academic.
- Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory*. New York: Oxford University Press.
- Chalmers, D. (2010). The singularity: A philosophical analysis. *Journal of Consciousness Studies*, *17*, 7–65.
- Corabi, J., & Schneider, S. (2012). The metaphysics of uploading. *Journal of Consciousness Studies*, *19*(7–8), 26–44.
- Corballis, M. (2009). The evolution and genetics of cerebral asymmetry. *Philosophical Transactions of the Royal Society B*, *364*(1519), 867–879.
- Coren, S., Ward, L., & Enns, J. (2004). *Sensation and perception* (6th ed.). New York: Wiley.
- Dainton, B. (2014). Supplement to temporal consciousness. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy (Spring 2014 Edition)*. <http://plato.stanford.edu/archives/spr2014/entries/consciousness-temporal/>
- Efron, R. (1970). The minimum duration of a perception. *Neurophysiologia*, *8*, 57–63.
- Eth, D., Foust, J., & Whale, B. (2013). The prospects of whole brain emulation within the next half-century. *Journal of Artificial General Intelligence*, *4*(3), 130–152.
- Galaburda, A., Rosen, G., & Sherman, G. (1990). Individual variability in cortical organization: Its relationship to brain laterality and implications to function. *Neuropsychologia*, *28*(6), 529–546.
- Gazzaniga, M. (1967). The split brain in man. *Scientific American*, *217*(2), 24–29.
- Gazzaniga, M. S., Bogen, J. E., & Sperry, R. W. (1962). Some functional effects of sectioning the cerebral commissures in man. *Proceedings of the National Academy of Sciences of the United States of America*, *48*, 1765–1769.
- Hayworth, K. (2010). *Killed by bad philosophy: Why brain preservation followed by mind uploading is a cure for death*. Essay published online at <http://brainpreservation.org/content/killed-bad-philosophy>
- Hayworth, K. (2012). *Electron imaging technology for whole brain neural circuit mapping*. *International Journal of Machine Consciousness*, *4*(1).
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, *32*, 127–136.
- Jackson, F. (1986). What Mary didn't know. *Journal of Philosophy*, *83*, 291–295.
- Kandel, E., Schwartz, J., & Jessell, T. (2000). *Principles of neural science* (4 edn.). New York: McGraw-Hill.
- Knott, G., Marchman, H., Wall, D., & Lich, B. (2008). Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. *The Journal of Neuroscience*, *28*(12), 2959–2964.
- Kurzweil, R. (1999). *The age of spiritual machines: When computers exceed human intelligence*. New York, NY: Penguin Books.
- Kurzweil, R. (2005). *The singularity is near: When humans transcend biology*. New York, NY: Penguin Press.
- Levin, J. (2013). Functionalism. In E. N. Zalta (Ed.) *The Stanford encyclopedia of philosophy (fall 2013 edition)*. <http://plato.stanford.edu/archives/fall2013/entries/functionalism/>
- Levy, J. (1977). The mammalian brain and the adaptive advantage of cerebral asymmetry. *Annals of the New York Academy of Sciences*, *299*, 264–272.
- Ludlow, Peter, Nagasawa, Yujin, & Stoljar, Daniel (Eds.). (2004). *There's something about Mary: Essays on phenomenal consciousness and Frank Jackson's knowledge argument*. Cambridge: MIT Press.
- MacNeilage, P., Rogers, L., & Vallortigara, G. (2009). Origins of the left and right brain. *Scientific American*, *301*, 60–67.
- Nagel, T. (1971). Brain bisection and the unity of consciousness. *Synthese*, *22*, 396–413.

- Nida-Rümelin, M. (2010). Qualia: The knowledge argument. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (summer 2010 edition). <http://plato.stanford.edu/archives/sum2010/entries/qualia-knowledge/>
- Olson, E. (2010). Personal identity. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (winter 2010 edition). <http://plato.stanford.edu/archives/win2010/entries/identity-personal/>
- Oncel, D., Demetriades, D., Gruen, P., et al. (2007). Brain lobectomy for severe head injuries is not a hopeless procedure. *Journal of Trauma*, 63(5), 1010–1013.
- Parfit, D. (1982). Personal identity and rationality. *Synthese*, 53, 227–241.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Clarendon Press.
- Pockett, S. (2003). How long is ‘now’? Phenomenology and the specious present. *Phenomenology and the Cognitive Sciences*, 2, 55–68.
- Pulsifer, M., Brandt, J., Salorio, C., et al. (2004). The cognitive outcome of hemispherectomy in 71 children. *Epilepsia*, 4(3), 243–254.
- Putnam, H. (1967). Psychological predicates. In W. H. Capitan & D. D. Merrill (Eds.), *Art, mind, and religion* (pp. 37–48). Pittsburgh: University of Pittsburgh Press.
- Putnam, H. (1988). *Representation and reality*. Cambridge, MA: MIT Press.
- Robinson, H. (2012). Dualism. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (winter 2012 edition). <http://plato.stanford.edu/archives/win2012/entries/dualism/>
- Ruhnau, E. (1995). Time gestalt and the observer. In Metzinger (Ed.), *Conscious experience*. Paderborn: Schöningh; Exeter: Imprint Academic.
- Sandberg, A. (2013). Feasibility of whole brain emulation. In Müller, V. C. (Ed.), *Theory and philosophy of artificial intelligence (SAPERE)* (pp. 251–64). Berlin: Springer.
- Sandberg, A., & Bostrom, N. (2008). *Whole brain emulation: A roadmap, technical report*. Future Institute, Oxford University. <http://www.fhi.ox.ac.uk/brain-emulation-roadmap-report.pdf>
- Shoemaker, S. (1984). Personal identity: A materialist’s account. *Shoemaker and swinburne, personal identity* (pp. 69–132). Oxford: Blackwell.
- Spiers, H., Maguire, E., & Burgess, N. (2001). Hippocampal amnesia. *Neurocase*, 7, 357–382.
- Springer, J., Binder, J., Hammeke, T., et al. (1999). Language dominance in neurologically normal and epilepsy subjects: A functional MRI study. *Brain*, 122(11), 2033–2046.
- Stanley, R. (1999). Qualia space. *Journal of Consciousness Studies*, 6(1), 49–60.
- Swinburne, R. (1984). Personal identity: A dualist theory. *Shoemaker and swinburne, personal identity* (pp. 1–66). Oxford: Blackwell.
- Unger, P. (1990). *Identity, consciousness, and value*. Oxford: Oxford University Press.
- Vallortigara, G., & Rogers, L. (2005). Survival with an asymmetrical brain: Advantages and disadvantages of cerebral lateralization. *Behavioral and Brain Sciences*, 28(4), 575–589.