

Personhood, Persistence, and Constitution

by

Daniel Casey

A thesis submitted to the Faculty of Graduate and Postdoctoral
Affairs in partial fulfillment of the requirements for the degree of

Master of Arts

in

Philosophy

Carleton University
Ottawa, Ontario

© 2017, Daniel Casey

Abstract:

In the following essay I argue that psychological continuity provides the most suitable account of a person's persistence. This is in accordance with Derek Parfit's claim that what's important in cases of survival is not personal identity, but psychological continuity. The essay focuses on a criticism raised by Marya Schechtman. She argues that Parfit's account doesn't sufficiently distinguish between the metaphysical and practical concerns of a person's persistence, which she takes to create significant problems. I respond to her criticisms to show that these problems are not as significant as she claims, and that by attempting to address them, her own view suffers. The difficulty with her view is she provides too broad a concept of personhood. This is the result of confusing the causal with the constitutive features of personhood. Once these features are distinguished, we can see why psychological continuity provides a better account of a person's persistence.

Acknowledgements:

I am greatly indebted to my supervisor, Professor Andrew Brook, for all his help with this essay. I'm very fortunate to have had such keen, clear guidance. It has been a privilege that I will carry with me.

I would like to thank all of the great professors I have had the pleasure of studying with at Carleton: Professor David Matheson, Professor Annie Larivée, Professor Gabriele Contessa, and (visiting) Professor Pragati Sahni. I would especially like to thank Professor Myrto Mylopoulos for her advice and assistance, and Professor Christine Koggel for her invaluable research seminar and all the arm twisting she did to get me here—I think most of my fellow graduate students from that research seminar have similar arm burns, all of whom I'm grateful to for input and such.

I would like to thank my siblings Kate, Patrick, and Richard for always thinking highly of me even when there wasn't much to think highly of; and I would like to thank Jess, and her wonderful son Chase for all of their love and support.

Most importantly, I would like to thank my Mother and Father, for everything.

Big tings.

Table of Contents

Abstract	ii
Acknowledgments	iii
Table of Contents	iv
Introduction	1
Chapter 1 – Personal Identity and Psychological Continuity	5
1.1 Personal Identity over Time.....	5
1.2 Psychological continuity.....	10
1.2.1 Locke’s Account.....	11
1.3 Revised Psychological Continuity.....	16
1.3.1 Revisions to Locke’s account.....	17
1.4 Identity is Not What Matters	21
1.4.1 The All-or-Nothing of Identity and Psychological Continuity	21
1.4.2 The One-to-One of Identity and Divided Psychological Continuity.....	24
Chapter 2 – Contemporary Objections	29
2.1 Williams’ Inverse Thought Experiment.....	30
2.1.1 Response to Williams.....	32
2.2 Olson’s Non-Existing Fetus Argument.....	35
2.2.1 Response to Olson	38
2.3 Schechtman’s Criticism.....	42
2.3.1 Metaphysical vs. Practical Concerns.....	43
2.3.2 The problem of multiplicity.....	48
2.3.3 Locke’s Psychological Continuity and The “Forensic Unit”.....	51

2.4 Summary.....	54
Chapter 3 – Schechtman’s Criticism of Parfit and The Extreme Claim.....	56
3.1 The Shift to Neo-Lockean Psychological Continuity.....	57
3.1.1 Does Schechtman Dismiss these Changes Too Quickly?.....	59
3.1.2 Reductionism.....	62
3.2 The Extreme Claim.....	64
3.2.1 Parfit’s “Extreme Claim”.....	64
3.2.2 Parfit’s Response.....	67
3.2.3 Self-interest Theory Argument.....	70
3.3 Schechtman’s Extreme Claim.....	70
3.4 Summary.....	77
Chapter 4 – Schechtman’s Account	79
4.1 Personal Life View.....	80
4.1.2 Literal Identity.....	80
4.1.3 PLV.....	83
4.1.4 Homeostatic Property Cluster.....	84
4.2 Literal Identity and HPC.....	86
4.2.1 Determinacy and Clusters.....	89
4.2.2 Homeostasis.....	90
4.2.3 Privileged properties.....	95
4.2.4 Reducibility.....	97
4.3 Broad Views of Personhood and HPC Concepts.....	99
4.3.1 The Functionality of Analytic Definitions.....	99

4.3.2 What Goes Around, Comes Around.....	100
4.3.3 Summary.....	102
Chapter 5 – Constitutive vs. Causal Personhood.....	103
5.1 The Storm Analogy.....	104
5.2 Extended Mind Debate.....	108
5.2.1 Setting up the Boundaries.....	111
5.3 How this relates to P.L.V.....	114
5.3.1 Persons and the Coupling-Constitution Fallacy.....	116
5.3.2 The Slippery Slope of Extending Personhood.....	120
5.3.3 Carving Nature at its Joints.....	122
5.4 The Biological Component.....	125
5.5 Concluding Remarks.....	126
Bibliography.....	128

Introduction

Personal identity is a topic that has provoked significant debate in several areas of philosophy. From the most basic standpoint, the questions of what we are as persons, and what makes it so that we persist over time as the same person we once were, even when we might have changed in so many ways, shape this debate. These questions attempt to determine the qualifications for a person to remain identical with his or herself over time. How is it that I am the same boy, who, at five years old, stepped on a wasps' nest and was mercilessly stung as he ran for the safety of his father's arms? I remember little of the incident myself. Most of what I remember comes from what I've been told. This could just as easily be a story about some other little boy that I'd heard so many times and began reconstructing as a memory of my own—except, I believe that it was me. But what makes that true? Otherwise, I have few memories from that period of my life, my body looks nothing like that boy's, I think, act, and behave completely differently. In fact, there is very little connecting me with that boy at all, and yet we are taken to be the very same person. These are the sorts of questions that motivate the debate of personal identity: what are, precisely, those strands responsible for the persistence of a person?

Traditionally, this discussion has focused on the question of personal identity over time, or diachronic personal identity: what are the necessary and sufficient criteria for a person to remain his or herself over a period of time? Since it deals with the relation of identity over time, this question is usually treated from a metaphysical perspective. What is it that defines a person, and how does one persist over time? Developing an account of personhood that satisfies these criteria produces a definition that appears significantly limited when compared with how we typically conceive of people. Many of the features

we associate with personhood—appearances, behaviors, biological functions, etc.—are merely accidental and fall outside the scope of this debate. What are important for this debate are those conditions that are necessary and sufficient to ensure the continuity of a person. Here, the account that is most often taken to meet these requirements has been psychological continuity. That is, for a person to persist over time requires that some form of psychological connection holds between their different temporal stages. This accords with the general intuition that it is our psychological self that distinguishes us as persons.

While this account has a strong intuitive appeal, it faces several significant criticisms. Some of these criticisms focus specifically on psychological continuity as the criterion for personal identity. These criticisms question whether psychological continuity does, in fact, constitute the necessary and sufficient condition for a person to persist over time. Perhaps there is a better-suited candidate. Other, more recent, criticisms focus on the actual debate, questioning whether it adequately captures the importance of the concept of personal identity. According to these critics, there is a profound difficulty in attempting to reconcile a metaphysical account of personhood with the sorts of practical concerns that should, ideally, apply to such an account. Here, the logical concept of identity over time, and the many ways that we conceive of persons in contemporary thought, seem at odds. Does a metaphysical account of personhood—the numerical identity of a person over time— address the real importance of personal identity? How does an account of psychological continuity represent these different—metaphysical and practical—concerns, and does it do so coherently? These criticisms take slightly different approaches, but share a common apprehension about psychological continuity. This essay

will focus on developing and responding to these criticisms, particularly a version developed by Marya Schechtman, to show that these do not actually create an issue for accounts of psychological continuity.

Chapter 1 will begin by presenting an account of personal identity as psychological continuity. I will look at Locke's theory, and respond to some of the criticisms raised against his view. I will then present a more contemporary account of psychological continuity theory given by Derek Parfit. He incorporates many of the revisions to Locke's earlier account within his own, more radical view. He claims that in cases of survival, it is psychological continuity, and not personal identity that is important. This is the view that I will be defending over the course of the essay.

Chapter 2 will look at some criticisms of psychological continuity theory in general. Bernard Williams criticizes the thought experiments that are used in support of psychological continuity. Erik Olson argues for an account of animalism (a version of biological continuity) which he thinks better addresses the question of personal identity. I will respond to these criticisms and begin to outline a criticism developed by Schechtman. She claims that contemporary accounts of personal identity do not sufficiently distinguish between the practical and metaphysical concerns of personal identity.

Chapter 3 will look at how Schechtman's criticism plays out with regard to psychological continuity in general, and specifically Parfit's version of a person's persistence. While Schechtman thinks that the tension between practical and metaphysical concerns is an issue for all accounts of personal identity, she sees it as problematic for accounts of psychological continuity, particularly Parfit's version. It is

the reductionist component of his view that she thinks commits him to the claim that “identity doesn’t matter” and leads to the further problem of what Parfit calls the extreme claim. I will go over her arguments, focusing on the extreme claim, and look at some responses to show that this claim doesn’t present the problem that she sees.

Chapter 4 will outline Schechtman’s own account of personal identity, which she refers to as literal, identity. Her account is called the Personal Life View. First I will outline her view, then I will look at some of the problems she runs into by trying to accommodate too broad of a concept of personhood. She uses Richard Boyd’s concept of a Homeostatic Property Cluster to try and reconcile her concepts of personhood and literal identity. This, however, creates some significant problems in her overall position.

Chapter 5 will look at a particularly prominent concern with Schechtman’s account. This is her confusion of the causal and the constitutive features of personhood. I argue that Schechtman adopts too broad a concept of personhood because she confuses the causal relations of personhood with the constitution of personhood. To show how this plays out, I reference the extended mind debate to draw similarities between extending cognition beyond the boundaries of the brain, and extending personhood beyond psychological continuity. Once this distinction is cleared up, we can see where Schechtman’s account is mistaken. I then briefly address the biological concept of personhood proposed by Schechtman, and finish the paper with some closing remarks.

Chapter 1 – Personal Identity and Psychological Continuity

The focus of this chapter will be to assess the personal identity debate. The following section will begin by briefly outlining the debate. The second section will present an account of psychological continuity, some of the earlier criticisms directed at this account, and how more contemporary versions of psychological continuity theory have responded to these criticisms.

1.1 Personal Identity Over Time

Personal identity over time is a complex issue. It connects several different concepts with the continuity of a person over time, each of which brings its own controversy to the debate. First, there is the question of personhood: how does one define a person, and what features are essential to such a definition? Then, there are also questions of identity and identity over time: what are the necessary and sufficient conditions for something to be identical with itself, and for that identity to hold across different points of time. When taken together, these concepts make for a complex debate: not only are they difficult to navigate, but each carries implications that extend to other issues. The concept of personhood, for example, carries moral implications. Any metaphysical account of personhood should be aware of its moral implications; on the other hand, developing an account based on moral considerations can be problematic in its own right—for starters, moral concepts frequently change, but what constitutes a person's identity over time should be consistent. While negotiating these types of concerns might seem complex, the question of personal identity over time provides a

basis for the way we understand ourselves and the world which we inhabit, and so, I think, warrants the perseverance.

The need to recognize the continuity of a person over time is at the core of our conceptual, and practical, frameworks. Our entire system of knowledge and understanding is predicated on the belief that we exist, as persons, and that that existence is temporally extended. This temporally extended existence is described by the concept of identity over time. Since we are persons, this becomes personal identity over time. By establishing this identity, we situate ourselves within a world made up of what we recognize as other persons, and determine the relations through which we coexist. These relations form the conceptual and ultimately practical concerns that affect our lives. From satisfying the most basic physical needs to developing complex psychological and social systems—things like concern for oneself, an intimate connection to one's past, concepts of ownership, reward, punishment, responsibility—all of these are central to the beliefs, behaviours, and institutions that govern our existence; and all of these are grounded in the continuity of personal identity.

It's important to be clear that the philosophical debate over personal identity is primarily concerned with what it takes for a person to persist over time, and not what it takes for a person to retain a sense of self over time. In this respect, the term 'personal identity' can be misleading. This confusion can be attributed to a difference between the psychological and the philosophical uses of the term. In contemporary speech, the psychological use of personal identity has become far more common. Psychologists refer to identity as the way in which we self-identify, as individuals within a larger group or community. This takes into consideration different levels, or capacities, in which we

exist, and interact with each other. According to a psychological model, personal identity is the level at which we distinguish ourselves from other people—this is based on personal attributes (appearance, physical characteristics, abilities), as well as relations (goals, attitudes, possibilities)¹. Typically, when we refer to personal identity, it is the psychological concept that we have in mind. This includes such notions as having a "sense of identity," undergoing an "identity crisis," or "identity politics." These concepts frequently feature in contemporary discourse; however, they do not refer to the concept of "personal identity" that is of primary importance to the philosophical discussion.

In philosophy, personal identity refers strictly to the conditions under which a person persists over time. To understand this, we can think of the necessary and sufficient conditions required for an object to remain identical with itself, and then substitute a person for that object. Because the philosophical term is concerned with identity in this sense, the debate is focused on the essential features of personhood—features without which the classification of personhood would cease to be applicable. This develops a much sparser account of personhood than we might expect, especially if we're considering personal identity in the psychological sense, i.e., the psychological level at which we distinguish ourselves from other people. There are likely many features that contribute to our understanding of what constitutes a person, or what shapes our psychological sense of identity, that would fail to qualify as essential features in the philosophical context. Consider things such as appearance, physical features, attitudes, personal relations—all of these shape the way one commonly thinks of themselves as a

¹ This is according to earlier models of psychological personal identity developed by Erik Erikson "Identity: Youth and Crisis" (1968), and "Identity and the Life Cycle: Selected Papers" (1959). These have been the foundation for contemporary notions of psychological personal identity.

person, but few, if any, are essential for personhood. Take, for example, something as simple as having a human body, with two arms, two legs, a torso, neck, and head. We might be inclined to say that the continuity of such a body is essential for the continuity of a person. But what about someone who has lost a limb, or had an organ replaced? Certainly, we would say that they are still the same person despite having parts of their body altered or exchanged. What, then, is essential about that body? Is having a human body sufficient to be considered a person? What about a corpse? We probably wouldn't want to consider a lifeless body to be a person. Maybe, then, sentience or consciousness is a more appropriate candidate. But what about interruptions in consciousness? Does a person cease to exist in the periods for which they are not conscious? What connects these conscious episodes? Once we begin to refine these suggestions, the intricacy of the issue becomes apparent. It is this bare-bones understanding of personhood that is relevant to the philosophical discussion of personal identity over time.

Another way this distinction is sometimes expressed is in terms of *qualitative* and *numerical* identity. Something is qualitatively identical with something else if they share the exact same features. On the other hand, something is numerically identical with something else if it is one and the same thing. These correspond, more or less, with psychological and philosophical identity: psychological identity is concerned with the qualitative features according to which we distinguish ourselves from each other, whereas philosophical identity is concerned with whether or not a person is, quantitatively, the same person that existed at an earlier time. Derek Parfit makes the distinction between qualitative and numerical identity clear using the example of two white billiard balls (1984. 201). Since they are identical in appearance and share all of the same qualitative

features, the two billiard balls are qualitatively identical. Because they are two distinct balls, they are not numerically identical; however, each ball is numerically identical with itself. If I were to paint one ball red, then it would no longer be qualitatively identical with the other ball, or with the white ball it was before it was painted red. It would, however, still be numerically identical with the white ball it was before being painted since it is still the same ball. In the same way, a person might undergo significant changes such that they are no longer the same person they once were in a qualitative sense—they might change their assigned gender, attitudes, beliefs; however, they would still be the same person in a numerical sense. It is numerical identity that is important for the philosophical discussion of personal identity.

Because this concept of identity is so central to our everyday life, and everyday life incorporates many concerns that are affected by both numerical and qualitative identity, these terms don't always separate as cleanly in practice as they do in theory. In some of the more complex cases, the distinctions become less obvious, and there are certainly examples where the two uses of identity overlap. When we begin to identify objects according to their qualitative features, for example, this distinction becomes less clear. If I were to ask you for the *red* ball, for example, you'd have to say there no longer is such a ball, since it has been painted white. There is still a ball, but not the red ball I am referring to. If you then became frustrated with the whole situation and brought a hammer down on the white ball, cracking it; and I were to say to you "fine, just give me any ball", you'd then have to tell me there no longer is a ball, since it has been cracked. There are now two pieces of a ball, which are still the same quantitative substance that was once the ball, however, sufficient qualitative changes have been brought about to that red ball to

say that it is no longer even a ball, let alone a red ball. For the case of persons, when we might have different qualitative concepts of what exactly constitutes a person, we can see how this distinction can become unclear.

It is, however, important to make the distinction between these terms to set out where they come together and where they should be separated. The general format of this debate adheres to these distinctions; and, typically, it is numerical identity that is important for philosophical discussion. Still, it's worth recognizing the limitations of these terms. For now, we will accept this distinction and focus on the earlier stages of the debate. Traditionally, the main options for the criteria of personal identity over time have been mind, body, and soul, or some combination thereof. And of these, psychological continuity has come to be considered the most viable.

1.2 Psychological Continuity

The most widely held philosophical view on personal identity is psychological continuity. This is the idea that for a person to remain the same person over a period of time, he or she must have some form of psychological continuity that connects his or herself at these different times—this is basically the idea of tracking psychological features backwards in a causal continuum. There are several ways of interpreting this continuity. The point, however, that these views share is that psychological continuity constitutes personal identity. Locke is often credited with being the earliest advocate of this view, and his account provides the foundation for most contemporary accounts of personal identity as psychological continuity.

1.2.1 Locke's Account

In chapter 27 of his *Essay Concerning Human Understanding*, Locke examines the concepts of identity, personhood, and personal identity over time. In this analysis, he famously distinguished the idea of a *person* from that of a *man*, describing the former as “a thinking intelligent being, that has reason and reflection, and can consider itself as itself, the same thinking thing, in different times and places” (Locke 1741, 286). The importance of this distinction is that *man* comes to be identified with the body, and *person* with the conscious being, capable of self-reflection.

Here, Locke takes consciousness to be the defining feature of personhood: “it does [consider itself] only by that consciousness which is inseparable from thinking, and, it seems to me, essential to it” (1741, 286). It is having this self-awareness, being able to conceive of oneself as oneself, that provides the basic constitution of a person. Locke focuses specifically on the cognitive sense of propriety that determines oneself to be the subject of one's own mental events: “When we see, hear, smell...or will anything, we know that we do so” (1741, 286). This sense of self-awareness accompanies all of our mental events— “[t]hus it is always as to our present sensations...and by this every one is to himself that which he calls self” (Locke 1741, 286)— and ultimately provides the foundation for our concept of self. “[I]n this alone [consciousness across time] consists personal identity” (Locke 1741, 286). It is this concept of self, which is understood to be consciousness, that Locke takes to constitute personal identity.

The importance of this claim can be emphasized by a distinction between consciousness and substance. According to Locke, personal identity does not amount to substantial identity: “whether we are the same thinking thing, i.e, the same substance or

no...concerns not personal identity at all: the question being, what makes the same person, and not whether it be the same identical substance" (1741, 287). Since different thinking things might constitute a single person, personal identity must not necessarily entail the identity of a substance. Rather, it is the same self-awareness that accompanies these different thinking things— "different substances, by the same consciousness (where they do partake in it) being united into one person" (Locke 1741, 287). For Locke, it is the sameness of consciousness (remembering now from the same point of view in which I was engaged then) that contains these, potentially discontinuous, substances, which is responsible for the identity of personhood. This establishes a concept of personal identity that can be understood independently of the human body since the criterion for personal identity is the continuity of consciousness, which doesn't necessarily pertain to the continuity of a single underlying substance.

While this criterion might seem narrow, Locke gives good reason for supporting his position. The principal motivation is implicit in the self-ascription of personhood. In order for one to consider oneself and others as persons, one must be capable of self-awareness. One's awareness of oneself as a being that is capable of conscious awareness, then, seems to be a necessary constituent of personal identity. If one could not conceive of oneself as oneself, then, for Locke, this would pose a problem for personal identity. This point is crucial:

For it being the same consciousness that makes a man be himself to himself,

personal identity depends on that only, whether it be annexed solely to one individual substance, or can be continued in a succession of several substances. For as far as any intelligent being can repeat the idea of any past action with the same consciousness it had of it at first, and with the same consciousness it has of any present action; so far it is the same personal self. (Lock 1741, 287)

Accordingly, consciousness, as self-awareness, is required for personal identity. The consciousness Locke describes incorporates this self-awareness with a form of memory connection, linking an awareness of past mental events with present mental events through the same conscious point of view. This reduces consciousness to that sense of propriety that accompanies all of one's own mental events—the unique relation that mental events bear to the subject that experiences them as its own. Thus it isn't the continuity of the mental events themselves, but the continuity of the psychological point of view that witnesses these mental events that constitutes personal identity. This is the basic premise for psychological continuity that has been so influential. It captures the general intuition that consciousness is fundamentally linked with personal identity, which, when considered along with the possible discontinuity of an underlying substance, makes a strong case for accepting psychological continuity as the sole criterion of personal identity over time. Despite this seemingly strong case, however, developing an account of psychological continuity that satisfies the concerns of personal identity isn't straightforward.

Although he is quite thorough on the subject, Locke isn't overly specific with exactly how he understands the sameness of consciousness. At times he seems to call for

a direct memory connection to past events— "as far as this consciousness can be extended backward to any past action or thought, so far reaches the identity of that person" (Locke 1741, 286-7). This claim has opened his theory to criticism since it seems to suggest that personal identity extends only as far as memories. If we take this memory connection to be sufficient but not necessary for personal identity, then there isn't a problem; however, if we take this memory connection to be both necessary and sufficient for personal identity, this is a different story. Thomas Reid argues, according to the latter interpretation of Locke, that a man might both be, and not be, the author of an action. He outlines a scenario where a brave soldier is flogged as a boy for stealing an apple, captures the enemy standard at his first battle, and is promoted to General later in life. Suppose the old General recalls standing proudly on the battlefield with the enemy standard, and the soldier on the battlefield recalls being beaten for stealing an apple; but the old General has no recollection of ever stealing that apple—if this Lockean account is taken to entail the necessary and sufficient conditions, then these would have to be different people (Reid 1785/2008, 114-5). This problem, sometimes referred to as the young officer paradox, results from the transitivity of identity: for an identity to hold over time, there must be an identity relation that connects each temporal part within that period of time. If the boy is identical to the soldier, and the soldier is identical to the general, then, logically, the boy must be identical to the general. This is simply the concept of numerical identity, which, if Locke's account is taken to pertain to the necessary and sufficient conditions of personal identity over time, it doesn't appear to meet.

Another early criticism of Locke's account relates to the potential circularity that is involved in defining personal identity by consciousness when it would seem that having conscious memory already implies having personal identity. This criticism was first raised by Joseph Butler, who expressed the problem as follows: "one should really think it self-evident, that consciousness of personal identity presupposes, and therefore cannot constitute, personal identity, any more than knowledge, in any other case, can constitute truth, which it presupposes" (1736/2008, 100). The point, here, is that since consciousness of personal identity across time already implies personal identity, the necessary and sufficient condition for personal identity cannot be consciousness. Otherwise, this leads to a circular regress: I am now at (t1) the same person I was then at (t2), because now at (t1) I have conscious memories of then at (t2). Those memories indicate that I was then at (t2) the same person I am now at (t1), because now at (t1) I have conscious memories of then at (t2) ... Conscious memories cannot disclose whether or not I am now the same person I was then because in order to have those conscious memories I must already be the same person that had those memories, which is circular. Butler's point is, if having memories already presupposes personal identity, then memories cannot be used to define personal identity.

One final criticism that is often raised against Locke's account is that the focus he places on memory is unnecessarily restrictive. Having memory serve as the primary determinant of consciousness, and subsequently, personal identity, limits certain cases, as illustrated in Reid's earlier example, that would intuitively be classified as cases of personal identity. There are many other psychological relations, including beliefs, desires, intentions and their actuations, that would be similarly suitable contents for psychological

continuity. Self-referential memory is an obvious candidate because it refers directly back to a previous time; however, it is not the only relevant psychological feature. This is taken to be another criticism of Locke's account.

Together, these make up the most significant criticisms of Locke's version of psychological continuity. More recent accounts of psychological continuity have responded to these criticisms in several ways. Derek Parfit provides an account that consolidates these responses to problems with Locke's earlier view into what might be considered a more radical version of psychological continuity, which will provide the main account of psychological continuity for this essay.

1.3 Revised Psychological Continuity

In an essay entitled "Personal Identity" (1971), Derek Parfit proposed a radically different way of addressing the question of personal identity over time from the standpoint of psychological continuity. According to his account, implicit in the issue of personal identity is the assumption that questions about personal identity can be given a determinate answer. When asked whether or not I will have continued to exist at a later date, I expect this question to have a determinate answer. I would assume the answer would have to be: either yes, I have continued to exist, and therefore my personal identity has been retained; or no, I haven't, and there is no longer a case of personal identity. This implies that the relation of personal identity is not one of degrees, but rather all-or-nothing. A second belief that is implicit with regards to personal identity, according to Parfit, is that having an answer to such a question—whether I have or have not continued to exist—is necessary in order to render significant practical judgments. In other words,

many of the important judgments that govern our daily lives—questions of survival, moral responsibility, and memory—all depend on there being a definitive answer to the first question of whether or not I will have continued to exist at such and such a point of time. Parfit challenges these assumptions.

Parfit's argument is, while the importance of these practical judgments happens to coincide with the concept of personal identity, their importance is not derived from the relation of identity, but rather what is represented by the relation of identity. This is because identity is one-to-one and all-or-nothing. The persistence of a person, however, involves neither of these. So, whatever it is that's going on in the persistence of a person, it is not identity. What matters (for judgments of moral responsibility, etc.) is whatever results in a person's persistence, not identity. And, what results in a person's persistence is psychological continuity of various kinds, memory being paramount among them. Once we relieve the debate from the confines of identity over time, he argues, what the question of personal identity is attempting to determine becomes much simpler.

In *Reasons and Persons*, Parfit expands this position into a more thorough account. This account addresses the concerns raised against the earlier Lockean model, incorporating them into a comprehensive, revised version of psychological continuity. We will first look at how his account resolves the problems with Locke's earlier account. Then, we will look at Parfit's specific claim that personal identity is not what matters.

1.3.1 Revisions to Locke's account

One of the first criticisms we saw of Locke's theory was the problem of transitivity. This problem was exemplified in the young officer's paradox given by Reid.

In order to resolve this issue, Parfit distinguishes the concepts of psychological *connectedness* and *continuity*: connectedness refers to a direct psychological connection, whereas continuity consists of an overlapping chain of psychological connections (1984, 206). For psychological continuity to hold, there doesn't have to be a direct memory connection with a past self; rather it is sufficient that there is an overlapping strand of memory connections linking each successive self. I may no longer hold any memories from when I was a child, so long as I hold some memories of a past self, that holds some memories of a past self...that holds some memories from when I was a child—this constitutes psychological continuity. If we accept this definition of psychological continuity over psychological connectedness, then we can see that the young officer's paradox no longer presents a problem for psychological continuity. Provided the old general has some memories that connect him to the young officer, and the young officer had some memories that connect him to the apple stealing boy, then the old general is transitively connected to the apple-stealing boy, even if he no longer holds any memories of ever stealing an apple.

To resolve the second criticism of Locke's account, Butler's claim that in order for a person to have memories of a past person, they must already be identical with that person, and memories, therefore, cannot define personal identity, Parfit introduces Sydney Shoemaker's concept of *quasi-memory* (Parfit 1984, 220). Quasi-memories are memories that might not entail our own experience by definition. The idea, here, is to allow for the possibility that we might have memory-like states that are not our own. Butler's criticism relies on the necessary connection between conscious memories and personal identity. If having conscious memories already requires that I am identical with

the person who experienced those memories, then having those memories cannot also explain why I am identical with the person who initially experienced them. This would be circular: I'm identical with a person because I have their memories; I have a person's memories because I'm identical with them. If, however, conscious memories didn't necessarily presuppose personal identity, then this would evade the circularity. If, for example, I could have memories that were not necessarily my own, then having memories would not necessarily imply that I am identical with the person who experienced the content of those memories. While personal identity could still be constituted by conscious memories, conscious memories would not require personal identity, and the relation between the two would not be circular.

What Shoemaker proposes is the logical possibility that memories do not entail personal identity (Shoemaker 1970, 281). Suppose, for example, it was possible to split a single human brain into two bodies and then reunite both parts of the brain within the same body. The possibility of this type of branching scenario would suggest that certain conscious memories do not necessarily constitute personal identity. This is because identity is one-to-one and splitting consciousness would create two separate entities which cannot, therefore, be identical. Being conscious of a memory when the brain is whole, that occurred while the brain was split, would be a case where having such a memory would not necessarily presuppose personal identity—otherwise, it would violate Leibniz's law and the transitivity of identity (Shoemaker 1970, 280). This is because the memory in question would have originated from one of two psychologically continuous beings, which is obviously not a one-to-one relation, and so cannot be described as a case

of personal identity. If we accept the possibility of these scenarios, then quasi-memories can be used to avoid Butler's charge.

According to this solution, ordinary memories are a kind of quasi-memory. Quasi-memory is simply a broader category than ordinary memory. Quasi-memories are *our* memories when they are causally connected to a certain experience that actually happened *to us*, in the *right kind* of way; and having these memories that are connected to experiences, supports the psychological continuity account of personal identity. Different accounts vary as to what exactly entails the *right kind* of cause. These range from the more conventional notion, where the right cause would be physiological (the continued existence of the brain), to more extreme interpretations, where the right cause could be *any* cause—these would include the types of sci-fi scenarios, where psychological continuity persists despite the destruction and reconstruction of the human brain/body. For Parfit, “in an account of what matters, the right kind of cause could be any cause” (1984, 215). Since Parfit's account, like Locke's, is committed to the distinction between psychological continuity (sameness of consciousness), and physiological continuity (sameness of life/substance), he is committed to the more extreme version of this claim, that memories connected by *any* cause are sufficient for maintaining personal identity over time.

In response to the third criticism of Locke's account, Parfit suggests that there isn't good enough reason to restrict this psychological connectedness to memories, and it should be expanded to include beliefs, desires, and the actuation of intentions (1984, 207). While memories provide an obvious connection to past stages of a person, they are not the only psychological connection that is relevant for cases of a person's persistence.

1.4 Identity is Not What Matters

The above amendments satisfy the three concerns directed at Locke's account of psychological continuity. Combining these responses with the traditional Lockean model, Parfit develops his revised version of psychological continuity. The most notable part of this revised psychological continuity is the claim we discussed earlier, that personal identity is not what matters. This is really the central component of Parfit's account. In order to fully appreciate this claim, it will be helpful to work through several sci-fi types of hypothetical scenarios. These thought experiments are meant to draw out two major points about identity that don't correspond with psychological continuity: (1) that identity is all-or nothing, and psychological continuity is a matter of degrees; and (2) that identity is one-to-one, and psychological continuity is capable of branching. Setting these points out makes the claim that in cases of a person's persistence, identity is not what matters, much clearer.

1.4.1 The All-or-Nothing of Identity and Psychological Continuity

In the first thought experiment we will look at what is known as the Psychological Spectrum Case. Parfit borrows from Williams' famous brain transplant scenario (which we will discuss further in the next chapter) to make a case for the indeterminacy of psychological continuity (Parfit 1984, 231). Briefly, if we accept a weak version of psychological continuity (personal identity as psychological continuity, but that psychological continuity requires the continuity of a brain), then in a brain-body swap scenario, we would agree that personal identity accompanies the brain. So, if person *X*

were to have her brain removed and transplanted into person *Y*'s body, then *X* would wake up in *Y*'s body. Now, suppose during this brain-body swapping procedure a portion of *X*'s brain is damaged. The doctor, however, undeterred, proceeds with the operation and is quite pleased with the results: the patient who wakes up in *Y*'s body possesses many of *X*'s memories. In fact, the number of *X*'s memories that have been successfully transferred into *Y*'s body turns out to be directly proportional to the amount of brain matter that remained undamaged throughout the procedure. Now, in accord with our earlier intuitions, we would probably agree with the doctor that the operation has been a success and that the patient waking up is indeed patient *X*.

It would seem we would be inclined to maintain this belief that the resulting person is *X*, regardless of the percent of brain that survived the transplant, provided she still has sufficient memories of her earlier self. That is to say, if *X* can still remember being *X*, then we would probably be inclined to say that she is indeed *X*. It shouldn't matter whether 90% of *X*'s brain survived the operation or 30% of it did. It would seem that she is still, numerically, the same person. Certainly, she hasn't become a new person; and, certainly, the resulting person is still, in fact, a person. Is there, however, a point at which the resulting person would no longer be person *X*?

Now, consider the same operation, except this time both *X*'s and *Y*'s brains are being preserved. Instead of putting *X*'s brain into an empty skull, the doctor has realized that in order for such a transplant to be successful and the body to accept a new brain, he must integrate a part of the body's original brain. Thus, he must splice a portion of *X*'s brain with *Y*'s brain, so that *Y*'s body will accept *X*'s brain. Once again, the operation is a success, and the doctor is pleased with the results. As before, the patient waking up

contains an amount of memory directly proportional to the amount of brain matter that was transplanted. Only this time, she possesses memories from both *X* and *Y*. This raises a complication for questions of personal identity: who is the patient waking up on the operating table? If only a small portion of *Y*'s brain were used to facilitate the transplant, then most likely the patient waking up would still be *X*, who just happens to have a few of *Y*'s memories as an unfortunate consequence of the operation. However, what if a significant portion of *Y*'s brain were used: say 50%? Is there a clear point at which the person would become someone else?

If we are committed to the claim that personal identity determines this relationship, since personal identity is all-or-nothing, there would have to be a determinate point at which the resulting person could no longer be considered *X* and would have to be considered *Y*—a threshold of memory connectivity after which personal identity must be considered lost. The trouble is determining where exactly such a threshold would exist. The obvious suggestion might be to say 50% memory connectivity is necessary in order to retain personal identity, anything less and a person is no longer the same person. This, however, is slightly troubling. On the one hand, it raises the question of what becomes of a person that possess less than 50% memory connectivity with a past self? If there is still a person and she remembers 49% of her past self, wouldn't we be inclined to still consider her to be the same person? Is that 1% of memory sufficient grounds to claim that a person either has, or has not, continued to exist? It seems trivial to say the difference between being and not being the same person amounts to 1% psychological connectivity. Conversely, if the person is still psychologically continuous with their past self, only we've decided that they no longer meet the

requirements of personal identity, then it becomes arbitrary. Either way, it doesn't appear that the question has a determinate, all-or-nothing answer.

It is what Parfit refers to as an empty question—a question where we can know all the facts without having a clear answer (1984, 232-3). In cases of empty questions, we might know all there is to know about a state of affairs without being able to render a decision that isn't arbitrary. We might decide that 50% is the cut-off of personal identity, but this would simply be an arbitrary benchmark. This scenario suggests that what matters in a person's persistence might not, in fact, be determinate, or all-or-nothing, as would be the case with identity. The second thought experiment we will look at brings up the possibility of division and branching in psychological continuity. If identity describes a one-to-one relation, and psychological continuity allows for the possibility of branching, then it doesn't seem that personal identity could account for a person's persistence as psychological continuity.

1.4.2 The One-to-One of Identity and Divided Psychological Continuity

As we saw in the Psychological Spectrum Case above, if we accept an account of personal identity as psychological continuity, then there appear to be cases where questions about personal identity would not have clear answers. This is because personal identity refers to a relation that is all-or-nothing, and psychological continuity comes in degrees. This next thought experiment, which deals with psychological division, points out that a person's persistence as psychological continuity might branch out, and so would not take the one-to-one form that is necessary for numerical identity.

To support this thought experiment, Parfit refers to medical procedures where severing the corpus callosum was used to reduce the severity of epileptic seizures but had the unintended consequence of creating two separate spheres of consciousness corresponding with the different brain hemispheres (1984, 245). Further support is the possibility that an individual might be equipped with identical left and right brain hemispheres (Parfit 1984, 246). Together, these ideas can be taken to support the possibility that an individual with identical brain hemispheres might be divided into two psychologically identical individuals. Parfit explores this scenario in the thought experiment *My Division* (1984, 254). Here, he describes a scenario where both of his brain hemispheres are exactly the same, and he is one of three identical triplets. In this scenario, his body is destroyed, but his brain remains intact, and each of his two siblings' brains are destroyed, but their bodies remain intact. Then, one of his surviving brain's hemispheres is transplanted into each of his brain-damaged brothers' bodies, resulting in two individuals psychologically continuous with himself.

According to Parfit, this generates four possible outcomes: (1) he doesn't survive, (2) he survives as one of two people, (3) he survives as the other, or (4) he survives as both (1984, 256). He decides the only sensible outcome is the 4th possibility: that he survives as both. If this is not the case, then “[h]ow could a double success be a failure?” (Parfit 1984, 256). If what matters for a person’s persistence is psychological continuity, and in cases of division psychological continuity can be split across multiple people, then what matters for a person’s persistence can be split across multiple people. This, he argues, supports the claim that what’s important in cases of survival cannot be a matter of identity, since identity necessitates a one-to-one relation and this is clearly a case where

what matters splits from one to two. If personal identity were to matter, it would have to be one of the other three options. If psychological continuity matters, it must be the fourth option. This thought experiment forces the decision—personal identity or psychological continuity (Parfit 1984, 278)—whereby psychological continuity is the obvious solution.

Parfit then takes this conclusion one step further and applies it to the branching case (1984, 200-1). Branching cases are similar to the case of division, except they focus specifically on psychological continuity without the brain requirements. Because the case of My Division still pertains to the brain, it doesn't represent an account of pure psychological continuity, but rather a weak version of psychological continuity, or a variant of biological continuity. In the hypothetical case of Teletransportation, then, where your psychological constitution is uploaded to a computer and recreated elsewhere (say on Mars) if your original self were to remain psychologically continuous on earth, this would be a case of double survival. In such a case, you would be psychologically continuous with both individuals, which, according to Parfit, is just as good as ordinary survival.

There are several features of Parfit's account that are worth noting to make the reasoning behind this conclusion a little clearer. First, his account of personal identity is reductionist, which means that personal identity just consists in particular facts. There is nothing over and above these facts that constitutes personal identity (Parfit 1984, 209-13). This is an important feature that is at the root of Schechtman's criticism of Parfit, which will come up in chapters 3 and 4. According to Parfit, however, any account of a person's persistence that isn't reductionist is committed to dualism. If there must be

something more to personal identity than all the facts, whatever that extra something must be, if it cannot be accounted for by facts, then it cannot be anything physical. It makes sense, then, to hold a reductionist account of personal identity. And if we hold a reductionist account of personal identity as psychological continuity, then the facts that constitute personal identity are psychological connections. Parfit refers to this as Relation R: "psychological connectedness and/or continuity, with the right kind of cause"; and, in an account of what matters, he adds the more controversial claim, that this could be any cause (1984, 215). The first part of the clause stresses psychological continuity and/or connectedness. This suggests that psychological continuity is sufficient for a person's persistence. It's worth noting that Parfit alters this requirement slightly depending on the conditions: in branching cases, for example, connectivity becomes more important, whereas, in the absence of branching, continuity is sufficient.² This second clause allows for Relation R to be maintained in the more far-fetched hypothetical scenarios such as Teletransportation. So for Parfit, it is Relation R and not personal identity that matters.

While most people would find this objectionable and are of the mindset that personal identity does in fact matter, Parfit argues that this attitude is mistaken. It is true that what we take to be personal identity matters to us. It matters just because it coincides with psychological continuity. This, however, does not describe an actual case of identity; it is merely an approximation of it. Since personal identity doesn't exist, it would be a mistake to accord the concept with importance in its own right—it is only important because something similar to identity happens to coincide with psychological continuity. This becomes clear when we accept a reductive account of personal identity as psychological continuity and apply it to the branching scenario. If personal identity

² For a deeper discussion of this relation, see Belzer (1996).

consists in facts, and those facts are psychological connections, then there can be nothing more to personal identity than holding these psychological connections or Relation R. In the case of Teletransportation, the individual would hold precisely those psychological connections, and she must be treated as being important in the same way that we consider personal identity to be important: in terms of psychological continuity. According to Parfit, "[T]his is not because Teletransportation is *about as good as* ordinary survival. It is because ordinary survival is about as bad as, or little better than, Teletransportation. *Ordinary survival is about as bad as being destroyed and Replicated.*" (1984, 280) When we dispense with the notion that personal identity requires something more than just the facts, we realize that there is no further fact to a person's persistence. Personal identity simply consists of a series of psychological connections, and these are what are of importance. In fact, personal identity doesn't matter, what matters is relation R. For the rest of this essay, cases that might be described as personal identity will instead be referred to as cases of a person's persistence. The term personal identity will still come up, though mostly in reference to theories or ideas that are committed to the identity component of a person's persistence.

Parfit's account is considered one of the more widely accepted versions of psychological continuity. In the next chapter, we will look at some criticisms that come up against this account.

Chapter 2 – Main Objections

Despite being one of the more widely accepted accounts of personal identity, psychological continuity in general, and Parfit's account of a person's persistence more specifically, are not without critics. This chapter will look at some criticisms that come up for these accounts. We will begin with Bernard Williams' criticism of the thought experiments used to support psychological continuity. According to Williams, these thought experiments are misleading and can just as easily be reconstructed to support versions of biological continuity. Williams points out how thought experiments in general are not always as reliable as philosophers credit them with being, and if we take these to be the primary motivation for psychological continuity, then perhaps our convictions are misplaced. The first section of this chapter will outline William's argument and assess why it shouldn't be taken as problematic for psychological continuity. The second section of this chapter will look at a criticism directed at psychological continuity from the perspective of biological continuity. Erik Olson's argument in support of animalism (a version of biological continuity) is the most recent example of this criticism. We will assess Olson's arguments to show why they don't actually create the problems he sees for psychological continuity. In the third section, we will look at a criticism developed by Marya Schechtman. Schechtman's criticism applies to accounts of personal identity in general, and focuses on a tension between the concept's practical and metaphysical commitments. She thinks this tension is a problem for any account of personal identity; trying to resolve this tension, however, creates further problems for accounts of psychological continuity, particularly Parfit's version. In this chapter we will focus on the

concerns Schechtman sees arising for accounts of personal identity in general, and the further problems she sees arising for accounts of psychological continuity. In the next chapter, chapter 3, we will look at how Schechtman thinks these concerns create a particularly difficult situation for Parfit's account of a person's persistence. For now, we'll begin by looking at Williams.

2.1 Williams' Inverse Thought Experiment

One of the earlier criticisms of psychological continuity within the contemporary debate is given by Bernard Williams. His criticism isn't specifically directed at psychological continuity, but at the credibility of the thought experiments that support psychological continuity. However, since these thought experiments are one of the main supports for psychological continuity, it is worth discussing. His criticism focuses on "bodily transfer" or "brain transplant" thought experiments³. These types of thought experiments are frequently used in the literature to elicit intuitions that support psychological continuity. While the specifics differ from case to case, they appeal to the same underlying reasoning. Basically, two individuals, *X* and *Y*, swap bodies such that the consciousness of *X* wakes up in the body of *Y*, and vice versa. Typically, when confronted with this body swap, our intuitions tell us that the person remains with their consciousness. Thus *X* wakes up in *Y*'s body, and *Y* wakes up in *X*'s body. This is meant to show that psychological continuity and not biological continuity is necessary and sufficient for personal identity.

³ (Locke 1741, 291); (Shoemaker 1959/2008, 121).

To make his case, Williams proposes a variation of this thought experiment where one of the two individuals will be tortured, and the other will be awarded a considerable sum of money—\$100,000 (1970/2008, 181). When asked to choose which individual should receive the torture and which should receive the money, knowing that they are about to undergo a brain transplant experiment that will leave their consciousness intact in the other individual's body, each selects the opposite body for the reward and their own body for the punishment. This decision supports our intuitions that personal identity corresponds with psychological, and not biological continuity.

Williams then asks us to conceive of another scenario (1970/2008, 185). Imagine you are being held captive and are told by your captor that tomorrow you will undergo excruciating torture; however, you won't remember having been told this, since just before being tortured, you will be made to forget it. Presumably, you would be terrified at the prospect of being tortured and knowing that you will forget having been told this in advance offers little consolation. Your captor then informs you that this removal of memory is part of a larger process, where all of your memories will be removed, and new memories will be implanted. Presumably, this too will do little to relieve your fear at the prospect of torture. Finally, your captor informs you that these new memories will correspond exactly with the memories of someone else. Still, the prospect of having your body tortured seems terrifying, regardless of what becomes of your memories.

This, however, is one side of the exact same experiment described earlier, with an emphasis placed on biological as opposed to psychological continuity. Williams argues that these types of thought experiments, which are a major motivation for psychological continuity, are inconclusive. They are structured in such a way as to elicit intuitions in

favor of psychological continuity—even the phrase “changing bodies” presupposed that personal identity is psychological (1970/2008, 198). Williams’ point is that the same thought experiments can be presented in such a way as to elicit a preference for biological continuity as necessary and sufficient for personal identity. In this second version of the thought experiment, it seems sensible that there would be a genuine consideration for our body that shouldn’t be present if psychological continuity were indeed necessary and sufficient for personal identity over time. Since these types of thought experiments are a major motivation for accounts of psychological continuity, Williams argues, there is no reason these accounts should be preferred over accounts of biological continuity.

2.1.2 Response to Williams

Williams' argument certainly does a good job of pointing out how our intuitions can be misled, although his argument is perhaps a better example of this itself than it is a critique of earlier thought experiments. He constructs a deliberately deceptive version in this thought experiment that depends more on the human capacity to project a concept of the self and emotions than its psychological predecessors. While the psychological versions of these experiments seem to elicit intuitions favoring psychological continuity based on straightforward scenarios, in order to elicit an intuition favoring biological continuity, Williams must significantly alter the presentation of his thought experiment. Also, instead of framing the issue as a direct question about personal identity, he uses the concept of fear to implicate a sense of personal identity, which isn't exactly the same thing—you might fear for someone or something without any indication that that person

or thing is you. We can see that this creates a much different case than those thought experiments that have traditionally been used in support of psychological continuity.

For Williams' thought experiment to successfully elicit intuitions in support of biological continuity, it must deliberately withhold presentation of the complete thought experiment. He begins, for example, simply focusing on personal continuity: "I am going to be tortured tomorrow" (1970/2008, 185). Already he's telling us this is him, "I," that is to be tortured, so this is taken to be a case of personal identity. He then adds the condition that he shall not remember being told this. This addition effectively weakens psychological continuity in the thought experiment; however, this is still a case of personal identity according to psychological continuity. The next amendment is that all of his memories will be erased, which terminates his psychological continuity, but leaves his biological continuity intact. Finally, we learn that new memories will be implanted, and "will exactly fit the past of another person now living" (1970/2008, 186). By delaying the complete details of the experiment, Williams effectively encourages a sense of fear of torture that was initially established in the first stage of the experiment when we were made to believe this was a straightforward case of torture. Had he simply presented the final outcome, it would not have the same effect. What's more, the question becomes drastically different if we learn that our memories will be implanted elsewhere or simply be terminated.

It would, I think, be quite unusual for someone not to fear the prospect of having their memories erased and their body tortured. Here, there seems to be a double cause for fear: one at the prospect of having your memories erased, and a second at having your body tortured. In the first case, having all of your memories erased, according to

psychological continuity, would be similar to death. It is unclear why, from the perspective of psychological continuity, we should be expected not to fear these results. In the second case, being worried about what will become of one's body doesn't necessarily support biological continuity. If you told me after my death, my body would be desecrated, I would likely be worried about this, regardless. It is quite possible that we, as persons, develop intimate connections with things—other persons, objects, animals—the welfare of which would cause us concern. Our body is a good candidate for something we share a particularly intimate connection with, and so is something we might, reasonably, have particular concern for. All this example suggests is that I have concern for my body, not that it is a matter of personal identity. Another point worth considering is the way Williams uses fear as an indication of personal identity.

Just because Williams' thought experiment elicits a sense of fear at the prospect of having your body tortured, does not mean you are persisting as the same person. It is possible that I might still possess a sense of fear at the thought of my body being tortured even if I recognized this as no longer being a case of me persisting as the same person. While I might no longer think of it as me that is being tortured, I might still feel a sense of propriety and responsibility for my body, and a sense of fear for the person that is about to be tortured. The human ability to project emotions onto past or future versions of oneself, or even onto other people is a well-documented phenomenon. The concept of empathy is essentially that: our ability to project emotions onto others, which allows us to develop a sense of social rapport. I'm not sure this is any different. The thought experiment basically consists in cultivating a sense of fear for oneself and projecting it onto a person that turns out to no longer be you. What is different, however, is comparing

this sort of intuition with the intuitions that a thought experiment without these mechanics would elicit. If for example, I were to simply ask whether you would wake up with your thoughts or with your body, this is much less controversial and, I think, a better way of framing the intuitions that support psychological continuity. While Williams' example is meant to show how these intuitions can be misled, I'm not sure it captures the intuitions in their most persuasive form, which is the real point of these thought experiments. It does, however, provide a caution about relying on thought experiments for supporting our theories. It's worth keeping this in mind when using these types of thought experiments for theoretical support.

2.2 Olson's Non-Existing Fetus Argument

Another, more contemporary, criticism of psychological continuity comes from Erik Olson. The motivation behind his criticism arises from a series of problematic implications that seem to result from an account of psychological continuity. One way he frames the issue is by raising the question: “[w]as I ever a fetus?” (Olson 1997, 1). According to Olson, psychological continuity theorists are forced to answer “no”. Attempting to resolve this, Olson claims, draws out several inadequacies to which the psychological continuity theorist would be committed. For personal identity over time, in order for a person to come from a fetus requires that they are numerically identical with that fetus—this is the transitivity of identity. However, for psychological continuity to hold, this couldn't be the case: the fetus is not conscious and therefore does not share psychologically continuity with the conscious person that it becomes. This raises the further problem of what, then, happens to the fetus: either it ceases to exist, or it becomes

a human animal, which would necessarily be distinct from the person, and yet simultaneously co-inhabit the same physical entity. According to Olson, neither of these suggestions is acceptable.

In his most recent book *What Are We*, Olson expands this line of criticism to any account of personal identity that is not a version of animalism. This includes most contemporary accounts such as the brain view and psychological continuity, as well as some of the more classic concepts such as Hume's bundle theory, accounts of souls, and nihilism about the self. While the specifics of his arguments for these different views vary on a case-to-case basis, his most central argument is similar to the fetus argument above.

He calls this *the thinking-animal problem* (Olson 2007, 29). The problem follows: when confronted with the question of what we are, the common response would be that we are human animals. There are just as many human animals on this planet as there are people, and it would seem that each person corresponds intimately with a human animal. If this is the case, then it would also seem that these human animals have the same mental properties: they share the same brains and nervous systems as do their person counterparts. Accepting this, then, commits you to the view that you are a human animal. Either that, or one of three alternatives outlined by Olson: (1) a human animal doesn't inhabit the same space as you, (2) a human animal does inhabit the same space as you but doesn't think the same way as you, or, (3) there is an animal that shares the same space as you and thinks exactly like you, but isn't you (2007, 30). Accepting the first alternative and denying the existence of human animals is senseless and commits you to a view that would likely deny the existence of many other things we take to exist (Olson 2007, 31).

Accepting the second possibility, according to Olson, commits you to one of two further possibilities: if we don't think like human animals, then either human animals don't think, or they don't think like us. Saying they don't think at all seems far-fetched, and saying they don't think the same way as we do, according to Olson, would have to fall along the lines of their being conscious, but not *self*-conscious. He maintains, however, that saying they don't think like us is just as senseless as saying they don't think at all (Olson 2007, 34-5). The third alternative—that human animals think just like we do—results in what Olson refers to as too many thinkers.

According to Olson, this too many thinker scenario presents three further challenges. If we are committed to the view that both a person and a human animal think the exact same thoughts and inhabit the same space, this creates three problems (2007, 35). (1) The overcrowding problem: that there are many more thinking beings than we understand there to be, since for every person there is also a thinking animal that inhabits the same physical space. (2) If the person and the thinking animal think the exact same thoughts, there is no way of distinguishing whether you are in fact the person or the human animal that is thinking. (3) If the thinking animal is exactly like the person in terms of thought, then shouldn't the thinking animal qualify for personhood, as well? Olson considers these three problems to present real challenges for any account of a person's persistence that isn't animalism. If we accept Olson's animalism, however, then none of these problems arise.

Overall, I don't find Olson's arguments persuasive. Most of them return to some variation of the too many thinkers scenario, which I don't think presents a real problem.

His view does, however, represent a contemporary criticism of psychological continuity theory, and so it is worth considering.

2.2.2 Response to Olson

As I mentioned above, I don't find Olson's arguments compelling. There are many responses, going all the way back to Locke's distinction between persons and human animals, available to the psychological continuity theorist. I will do my best to work through Olson's arguments in what I take to be a sensible manner.

Consider the first suggestion: was I ever a fetus? The answer to this question depends on the context in which I am considering myself. Here, there are two possible suggestions: either yes, I was a fetus, or no I wasn't. And these correspond to different accounts of personal identity: either I am a human animal, or I am a person that is not necessarily a human animal. Olson's claim that I was necessarily a fetus conflates the concepts of human and person. In response, you might simply say, as a human animal I was a fetus, but as a person, I was not. That is because a human animal grows from a fetus, and a person is a matter of psychological connections which a fetus is incapable of possessing. Think about this analogous question: was this oak tree ever an acorn? You might say that yes, that oak tree was an acorn since it is the same biological organism that once grew out of an acorn. Or, you might also say no, an oak tree is an oak tree and an acorn is an acorn. An oak tree couldn't be an acorn any more than an acorn could be an oak tree since one is a tree and the other a nut. In both cases, you would be correct. That is because we are talking about two different criteria of identity: in one case we are talking specifically about the identity criteria of trees and nuts, and in the other we are

talking about the persistence of a plant through different stages of growth. Once we clarify these different concepts, I don't think there is a problem. In fact, even Olson's own view would fall victim to this sort of an argument: if you were to develop identity conditions for an adult male, for example, then you might generate the equally counterintuitive position that an adult male was never a fetus since each is defined by different criteria. This, however, ceases to be an issue once you distinguish the different concepts of identity and persistence you are working with.

There is a very simple distinction, here, that Olson disregards, and that is the difference between the human animal and the person. To make this clear, we simply need to conceive of cases where human animals might not be considered persons, or cases where persons might persist beyond the extent of a human animal's lifespan. Olson provides an example of the first case for us: a fetus is a human animal that isn't a person. Another example that comes up in the literature is a de-cerebrated corpse with enough of the brain-spinal cord intact to maintain bodily function. This would be a case that would qualify as a human animal but would likely fail to be considered a person. It's worth keeping in mind, here, that the strength of these examples depends on whether or not one accepts the concept of personhood, and whether one considers it to be a matter of psychological continuity. If you don't accept that the concept of personhood has meaning beyond the human animal, or you have different identity conditions for personhood than psychological continuity, then either you consider the term personhood to be redundant, or you will have to come up with alternative examples that fit your criterion for personhood. I would maintain, however, that there is a genuine distinction between humans and persons, and that, for most people, the term personhood has real meaning—

for those who disagree, the next paragraph might present a more appealing argument. On the second side of this suggestion, while there are no cases of persons persisting beyond human animals to date, the literature is abundant with hypothetical brain-body swapping scenarios, cloning, and the recreation of psychological continuity. If we accept the possibility of any of these scenarios, which are becoming less and less far-fetched as technology advances, then we can likely see cases where personhood and the human animal come apart. In fact, if we don't appreciate this distinction, we are committed to the claim that all persons are human animals. But if we can appreciate this distinction at all, then we must recognize that these are different concepts and so must be treated differently.

Another way to go about addressing Olson's problem of too many thinkers is by accepting his second option. (2) A human animal does inhabit the same space as you but doesn't think the same way as you. Olson raises and disregards the suggestions that the difference between the thinking of a human animal and the thinking of a person might be the difference between consciousness and self-consciousness: a human-animal thinks in much the same way as a person, but a person possess the self-consciousness that distinguishes it from a human animal. He doesn't go into why he takes this to be no better than a human animal not thinking at all: "If human animals were incapable of having the same sorts of thoughts that we have, or if they could not think in the sense that we do, that would be just as surprising as it would be if they could not think at all" (2007, 35). Presumably, his reasoning would be something like this: if the human animal thinks in much the same way as the person, and the person thinks in such a way that he or she is

self-conscious, then why wouldn't the human animal be self-conscious as well? The problem with this approach is it refuses to permit the possibility of overlap.

It is possible for the human animal and the person to inhabit the same body. You might say, for example, that once the human animal becomes self-conscious, it becomes a person; or, you might say that the person is a product of self-conscious thought, which need not occur in a human animal, but simply results from psychological continuity occurring in whatever capacity it happens to occur. Olson denies this sort of overlap since it contributes to what he takes to be the thinking-animal problem, but this doesn't have to be the case. There are plenty of examples of a substance constituting different objects: a lump of glass, for example, could be blown into a vase. It would then qualify, simultaneously, as both a lump of glass and a vase. Not just any lump of glass would qualify as a vase, however. It takes certain properties to satisfy the criteria of a vase that wouldn't be present in just any lump of glass. Similarly, not all vases need to be made of glass—you might have a metal or plastic vase, for example. In the case of persons, the person and animal can be explained as different objects constituted by the same tissue. That tissue can satisfy the criteria for one without necessarily satisfying the criteria for the other, just like a lump of glass could be a lump of glass without satisfying the criteria for a vase. If we understand that two objects can be simultaneously constituted by the same substance, then we should have no problem accepting that a human animal and a person can exist within the same physical body.

Olson's arguments are some of the more well-known criticisms directed at accounts of psychological continuity from the biological standpoint. I take the responses I have provided here to satisfy the issue. There is much more that could be said, but for the

scope of this essay, this should suffice. The next criticism we will look at takes issue with accounts of personal identity in general. Understanding and responding to this criticism will be the main focus of the rest of this essay.

2.3 Schechtman's Criticism

In *Staying Alive*, Marya Schechtman develops a criticism that focuses on what she perceives to be an inadequacy in contemporary accounts of personal identity. To formulate this criticism, she distinguishes two sets of concerns that relate to the concept of personal identity. These are the metaphysical concerns and the practical concerns. Dividing up the concept this way sorts out some of the motivations in this debate. This division brings up issues that might not be so apparent when treating the debate from a single perspective. Schechtman thinks this division creates a tension in all accounts of personal identity; however, she is specifically concerned with how it affects psychological continuity. Schechtman argues that accounts of psychological continuity are inadequate in their attempts to resolve these concerns. Her conclusion is that psychological continuity provides too narrow a criterion for personal identity. She uses this as the foundation on which to build her own account, a more inclusive account of personal identity called the personal life view, which we will look at in more detail in chapter 4. First, let's look at the problems she outlines for psychological continuity.

In order to set up her criticism, Schechtman divides the issue of personal identity into two sets of concerns: the practical and the metaphysical. She further attaches several other terms to the debate to help structure her criticism and, eventually, her theory.

Before we get started on her actual criticism, it will be helpful to work through these terms.

2.3.1 Metaphysical vs. Practical Concerns

When we address the question of personal identity over time strictly as a question of identity—what are the necessary and sufficient conditions for a person to remain numerically identical over a period of time—we are effectively conducting a metaphysical inquiry into the nature of personal continuity. We are attempting to determine exactly what constitutes a person, and what makes a person continue as the same, numerically identical, person over time. If we accept this as a case of numerical identity, these questions have clear, determinate answers, which, essentially, provide the foundation for what it is that we understand ourselves to be. Ideally, if such an account were correct it would provide a consistent basis to which more specific, practical concerns might be directed. These concerns relate to our actions, behaviors, and intentions, and how they are received and influenced by the social structures we inhabit. In short, they correspond with our day-to-day concerns, and typically fall under the classification of moral judgments. The relationship between metaphysical and practical concerns should be relatively straightforward—after all, they are concerned with the same subject. However, given the extensive network of practical concerns that influence our daily lives, these issues don't always match up as neatly as theory dictates—or at least that's what proponents of this argument would claim. And, when they don't line up, this raises the issue as to which concerns should figure primarily in our account of personal identity.

Consider the following three cases as a way of understanding how the metaphysical and practical concepts of personal identity might come apart.

In the first case, a young boy is born into poverty. Growing up, he is constantly hungry and uncomfortable. Fortunately, this situation improves, and he eventually finds himself well-fed and in a stable home. As he gets older, however, although his situation has improved, he still remembers the feeling of starving and being without a place to sleep and formulates a very clear intention that he will never let himself become that destitute again. Thus, he commits himself to working hard and constantly pushes himself towards greater successes. Many years later, that same boy, who has now become a wealthy man, continues to work and push himself just as hard as when he was a boy, always carrying those vivid memories in the back of his mind. Even though he now faces no danger of poverty, those memories continue to influence his daily actions.

Now, consider the same boy. In this second case, however, just when his situation begins to improve, and he has a roof over his head, things turn for the worse. Fearing a return to poverty, the boy decides to take matters upon himself. Driven by a fear of starvation, he devises a plan to steal money from his schoolmates. He succeeds in stealing a considerable sum before he is eventually apprehended. He is convicted and sentenced to a juvenile detention center. Upon his release, he decides to change tack: he channels his fear of poverty into productivity and, as in the previous case, commits himself to working hard and pushing himself towards greater successes. Because he is a minor, this small matter of fraud is expunged from his record, and he goes on to become a very successful

man. Despite this, he always remembers being that starving child: a memory that constantly drives his success.

Consider one final scenario. In this third case, similar to the first, the boy growing up in poverty commits himself to working hard and constantly pushes himself to succeed. As an established businessman, however, things don't go so smoothly. The man makes several risky investments that leave him bankrupt. Once again, fearing a return to poverty the man becomes desperate. He convinces several friends and business associates to invest in what amounts to an elaborate fraud. In this case, as before, he is apprehended, convicted and sentenced to prison. Upon his release, he commits himself to work hard and succeed; however, this proves to be more difficult than in the previous case. He now has a criminal record, and very few places will hire him. Despite his commitments to work hard and succeed, none of his former associates will work with him, and no governing financial body will allow him to manage or trade funds. To top it off, any money he makes goes towards paying off the people he defrauded. So, he spends the rest of his life struggling with the poverty he so desperately wanted to avoid.

These three scenarios illustrate some of the concerns that relate to the concept of personal identity. If we accept an account of psychological continuity, then it's clear that these cases would all qualify as examples of personal identity. In each of the three cases, the man vividly recollects his youth, and the conscious memories he experiences go on to influence him through to a much later stage of life. Thus we have several distinct strands of psychological continuity—memory, belief, and the actuation of intention—that connect the poor boy with his older self. If we consider these examples, however, from a

practical standpoint, we don't reach such a uniform conclusion. The first case is perhaps the clearest example of a case where the metaphysical and practical concerns line up. In this scenario, the conscious memories and intentions that support psychological continuity, correspond directly with the prudential concerns that influence the boy's attitudes and behaviors over the course of his life. In other words, we might say that he is the wealthy man who constantly strives for success because he was the poor boy who almost starved. Thus, his practical life's trajectory neatly mirrors those strands that support his psychological continuity.

In the second case, however, we begin to see the metaphysical and practical concerns come apart. Here, the practical concerns that relate to personal identity are slightly different. In the first case, they corresponded with the boy's prudential concern for himself; in this second case, while certainly, his prudential concerns are still relevant, there is now also a matter of legal concerns. It is very clear that in both cases, the boy is numerically the same person. From a moral perspective, though, he seems to have undergone significant changes. The boy who defrauds his classmates is punished for his actions. This punishment expires, the charges are expunged from his record, and he goes on to become a successful businessman. Legally, we would be inclined to say that the successful businessman is not the same legal entity who had committed fraud. His punishment has been served, the crime has been removed from his record, and he is, effectively, a new legal entity. Thus within a single, numerical, personal identity there would appear to be different practical cases of identity.

The third case adds a little more to the domain of practical concern. As in the first two cases, there are still the matters of prudential and legal concern. Now, however, there

is also a matter of institutional concern, relating to the business man's exclusions from financial firms; there is the social concern, that arises from his friends' unwillingness to trust him; and there is financial concern, in terms of the restitution he must pay. In none of these cases is there a doubt that the man is numerically identical throughout the example. There are, however, multiple dimensions of practical concerns that relate to his numerical identity. After being convicted of fraud, the man is no longer considered to have the same moral identity as the hard working business man he was before committing the crime. This is clear from the way he is held in legal, institutional, social, and financial disregard. Unlike in the second case, where the boy's sentence is complete and his criminal record expunged, as a grown man these concerns are not so easily rectified. After serving time in prison and reintegrating within society, the man might succeed in restoring the legal status of his identity. His exclusion from participating in financial enterprises will remain throughout the duration of his biological life. The mistrust of his social spheres is entirely contingent on the individual relations of which they are comprised. And his financial burden is a function of the rate at which he pays it off. This distinction seems all the more poignant if we consider that the boy in the second case and the man in the third are guilty of committing very similar crimes. What is the difference that warrants such a drastic variance in practical concern? In other words, what is the relationship between our metaphysical and practical identity?

As we've seen, from the perspective of psychological continuity, all three cases present clear-cut examples of numerical identity. The practical concerns, however, do not appear to be quite as uniform. There are significant variances in the practical statuses of a single case of numerical identity. Depending on how we qualify moral entities, these

changes could be understood to imply separate moral entities. In fact, many of these different sets of practical concerns reflect the various meanings that being a person takes on in everyday life, and the different concerns that motivate questions about whether or not a person is the same person in any one of these colloquial senses. Are you the same person as the boy who stole \$5 from his mother's purse? Is that seven-year-old girl a person whose vote we can count on in the upcoming election? Is that man the same person whose DNA was found at the scene of the crime? When we refer to someone as a person, or as being the same person, in everyday speech, there are a variety of different things we could have in mind. We might be considering whether they have the same personality, whether they are morally responsible for the same actions, whether they should be treated as the same legal entity, or whether they are, in fact, the same biological organism. These are all different ways to conceive of personhood, and they correspond with the different sorts of practical concerns that are important for personal identity. If we were to classify these practical statuses as distinct moral entities, such as persons, then it would be the case that a single numerical personal identity could contain multiple practical personal identities. For Schechtman, this creates a real problem for accounts of personal identity.

2.3.2 The Problem of Multiplicity

When we divide the concept of personal identity according to practical and metaphysical concerns, it appears that practical concepts of identity go beyond numerical identity. In both the second and third cases described above, changes in the boy/man's practical status are not the result of changes in numerical identity, but rather result from

qualitative changes. While there is no question that numerical identity provides a locus for these practical concerns, the practical concerns extend beyond numerical identity and implicate a sort of qualitative character. Here, then, we can see where the issue of practical concern separates from numerical identity. Since our practical concerns relating to personal identity implicate identity in a qualitative sense, and many of the features that constitute the qualitative character of an individual are contingent, it becomes difficult to develop an account of the necessary and sufficient conditions for personal identity in a practical sense than doing so for an account of numerical identity. This difficulty is responsible for what Schechtman refers to as the *problem of multiplicity*: "If we include all the other relations that make up our form of life, the project of finding a single relation that defines the appropriate locus of our practical concerns appears, at least initially, hopeless" (2014, 81). The problem here seems to be that if we want an account of numerical identity, we won't be able to accommodate the practical concerns that relate to personal identity.

According to Schechtman, there are several different ways to interpret the relationship between metaphysical and practical concerns regarding the question of personal identity. These are *the independence model*, *the dependence model*, and *the coincidence model*. The first way, the *independence model*, is to treat them as distinct concerns. Because of the plethora of practical concerns that relate to personal identity, according to this model, it is best to treat the metaphysical and practical concerns of personal identity separately. Exactly how independent of each other these concerns should be taken depends on who we consult. Olson's animalism, for example, is a view that proposes the question of personal identity, when taken as a metaphysical question, is

best explained by the human animal. He maintains that practical questions obscure the issue and it is a matter of biological continuity that determines personal identity (Schechtman 2014,49). For her part, Schechtman disagrees with this. She argues that any account of personal identity requires what she calls a “forensic unit” (which will be explained shortly) to provide a suitable target for moral consideration, effectively uniting practical concerns with numerical identity (2014, 56-7). She argues that any account of personal identity must relate practical and metaphysical (which she calls literal) concerns.

This is the model Schechtman favors for personal identity. She calls it the *dependence model*. On this account, personal identity and practical concerns must be codependent. Personal identity is conceptually dependent on practical considerations since personal identity must provide a suitable target at which to direct practical judgments; and practical judgments depend on personal identity because personal identity must be present for practical judgments to be made (Schechtman 2014, 41). According to Schechtman, accounts of psychological continuity fall under the third type of relation: the *coincidence model*. On this account, practical and metaphysical concerns coincide: “the limits of the person must coincide with the limits of particular practical judgments” (Schechtman 2014, 41). That is to say that the metaphysical status of personal identity is determined by the practical considerations of personhood. Schechtman considers this to present a serious issue for accounts of psychological continuity, a problem which she traces to the shift from Locke’s psychological continuity to neo-Lockean psychological continuity, which will be examined in the following section.

2.3.3 Locke's Psychological Continuity and The "Forensic Unit"

To make her case for the distinction between coincidence and dependence models of personal identity, Schechtman proposes an alternative reading of Locke. On this view, Locke's account of personal identity can be interpreted to suggest something quite different from how it has come to be understood by contemporary proponents of psychological continuity. This alternative reading takes its departure from the two renowned definitions Locke provides for a person. The first, which we looked at earlier, is "a person is a thinking, intelligent being...that can consider itself as itself, the same thinking thing in different places and at different times" (Locke 1741, 286). The second comes slightly later in the text, where Locke defines a person as "a Forensic Term appropriating Action and their Merit; and so belongs only to intelligent Agents capable of a Law, and Happiness and Misery" (Locke 1741, 296-7). Between these two definitions, it is clear that: (1) consciousness constitutes personal identity; and, (2) that practical concerns are integral to Locke's understanding of a person. According to Schechtman, when Locke distinguishes sameness of substance from sameness of consciousness, arguing for the latter as the appropriate designation of personal identity, this argument suggests that the sameness of consciousness is necessary for ascribing the sort of moral responsibility that inherently defines a person. Basically: a person is something that can be held accountable for its actions; and so, it must be capable of following law and understanding the relation between actions and consequences. For this to be possible, a person must have consciousness (Schechtman 2014, 12). The specific relationship between sameness of consciousness and the ascription of practical concerns is where Schechtman proposes deviating from the traditional reading of Locke.

According to the traditional interpretation, a person is delineated according to the limits of ascribing moral responsibility. That is, a person extends only as far as their consciousness can reconcile an action with its consequences. According to Locke, a person cannot be held accountable for actions of which they are not conscious: “therefore whatever past actions it cannot reconcile or appropriate to that present self by consciousness, it can be no more concerned in, than if they had never been done; and to receive...reward or punishment, on the account of such action, is all but to be made happy or miserable in its first being” (Locke 1741, 297). On this reading, the limits of a person must coincide with the limits of moral responsibility. This results in the coincidence model. A second way of interpreting this relation between personal identity and practical concerns is by suggesting that they depend on one another. This is Schechtman’s dependence model. To see the distinction, we must appreciate some subtle differences.

The first thing Schechtman draws our attention to is the difference between attribution and responsibility. When we make a moral judgment, this can be broken into two steps. First, there is the matter of attribution: who is literally the cause of an action. Then, there is a question of responsibility: can the person to whom we've literally attributed the action be held morally accountable? The first part of this judgment, attribution, directs our attention towards a suitable moral target and the second part, responsibility, determines whether such a target can justifiably be held accountable.

These two steps of determining accountability bear distinct relationships with how we conceive of personal identity. The first step of this judgment is concerned with locating a suitable target to which questions of moral judgment can be directed. Schechtman calls this target a “forensic unit”—this is where a person is conceived of as a

suitable target for moral consideration. The second step, determining whether moral responsibility is warranted, relates with what Schechtman calls the “moral self” (2014, 14-5). This is where a person is defined by the actions for which they are rightly accountable. Thus we can see two very different accounts of person emerging from this distinction. In the first case, a person is a forensic unit—a suitable target for moral concern (whether or not this concern is justifiable is a separate question). In the second case, a person is a “moral self” and defined specifically by what they are accountable for.

To clarify this distinction, imagine that a young child smashes a porcelain vase. While the child is clearly the forensic unit that caused the vase to smash, it isn't clear whether it is morally responsible for its action. Upon further investigation, one might decide that it is, in fact, the negligent babysitter who is responsible for the vase being smashed. In this case, the child is the forensic unit—it provides a suitable target to begin our moral investigation. It is not, however, a moral self. We cannot hold the child responsible for its action since it cannot conceive of the consequences and their moral implications. The babysitter, however, is responsible for the child, and so, indirectly, is responsible for breaking the vase. Attributing the literal action to the child gives us the direction for where we might actually place moral accountability. Schechtman admits that when we lay it out like this, the "forensic unit" appears to coincide with the human animal, and the "moral self" with psychological continuity; and this is how it has typically been understood. Schechtman's alternative interpretation is to suggest that the first part of determining moral accountability, the forensic unit, should pertain directly to psychological continuity, and not the human animal. On this reading, identity relations are not "coextensive" with moral judgments, as would be the case in a coincidence model

(Schechtman 2014, 16). Instead, identity simply provides a suitable target for where we might direct such moral judgments. In this way, personal identity doesn't automatically implicate moral accountability for those actions of which we are conscious, but must be in place before questions of moral accountability can justifiably be raised. It is a failure to observe this distinction which results in what Schechtman considers to be a decisive shortcoming in contemporary accounts of psychological continuity.

2.4 Summing Up

In this chapter we looked at some of the criticisms that face continuity accounts of personal identity. We began with Williams' criticism of the thought experiments that are frequently used in support of psychological continuity. Here, we saw that, although Williams' example doesn't rule out our intuition in support of psychological continuity, it does cast some light on the ultimately inconclusive nature of thought experiments in general. We then looked at Olson's arguments against psychological continuity. These arguments fail to recognize an important distinction in the question of personal identity between the human animal and the person. The too many thinkers problem, in particular, disregards the possibility that the same thing might simultaneously fulfill different concepts, and, thus, two things might overlap. In the third section, we began looking at Schechtman's criticism of accounts of personal identity in general. She sees a tension between the practical and metaphysical concerns of personal identity. This tension is a problem for all accounts of personal identity. Attempting to resolve this tension, however, creates further problems for accounts of psychological continuity. This problem is apparent in the traditional reading of Locke; however, Schechtman thinks it becomes

most serious in Parfit's revised version of psychological continuity. She sees this presenting a serious issue that ultimately results in what Parfit calls *the extreme claim*, an argument that Schechtman thinks is fatal for Parfit's version of psychological continuity. The next chapter will look at how Schechtman sees this problem arising for psychological continuity, Parfit's own interpretation of the extreme claim, and why it shouldn't be taken to be as problematic as Schechtman claims.

Chapter 3 – Schechtman’s Criticism of Parfit and The Extreme Claim

In the last Chapter, we began to look at Schechtman’s criticism of accounts of personal identity. That criticism depends on a distinction between the practical and metaphysical commitments found in questions of a person’s persistence. Schechtman thinks this tension creates a problem for all accounts of personal identity, but that this problem is particularly difficult for accounts of psychological continuity. In this chapter we will be examining the shift from Locke’s psychological continuity view to neo-Lockean psychological continuity, particularly Parfit’s version of psychological continuity. This shift is where Schechtman thinks the criticism she raises becomes most apparent. Basically she sees the problem existing in the traditional reading of Locke, but takes her alternative reading of Locke to offer a way out that is no longer available to later, revised versions. Because they are founded on what she takes to be a problematic reading of Locke, Schechtman argues that contemporary accounts of psychological continuity get saddled with this problem, which ends up leading to what Parfit has referred to as *the extreme claim*—the argument that without numerical identity in the strict sense, there isn’t good reason for a moral agent to have future-directed self-concern.

We will begin this chapter by looking at how Schechtman sees this problem developing for accounts of psychological continuity. The first section will look at the issues that arise in a shift from Locke’s account to more contemporary revisions of psychological continuity. We will then look at why some of these shifts aren’t actually as problematic as Schechtman interprets them to be. The main problem, however, that Schechtman sees resulting from these changes is that they lead to Parfit’s extreme claim.

The second section will outline how Parfit expresses this problem, and the response he develops for it. The third section of this chapter will look at Schechtman's version of the extreme claim to determine why she feels Parfit's response is inadequate. We will then look at how this worry can be resolved.

3.2 The Shift to Neo-Lockean Psychological Continuity

The problem for contemporary accounts of psychological continuity, according to Schechtman, is they that fail to clarify the distinction between a person's persistence as a forensic unit, and a person's persistence as a moral self. There have been significant changes made to Locke's original theory, but without making this distinction clear, the resulting accounts have developed a tendency to conflate these two aspects of a person's persistence. In this respect, it is as if contemporary accounts of psychological continuity are attempting to simultaneously define both the forensic unit and the moral self. A "double duty" that, Schechtman argues, places "very strong demands on the view, demands which...it is ultimately unable to meet" (2014, 22). To show this, she proceeds to work through the main changes that have been implemented in theories of psychological continuity, pointing out where they seem to confuse these two terms.

One change from Locke's memory theory to neo-Lockean psychological continuity that Schechtman points out is the shift from memory connections to psychological continuity. On Locke's account, consciousness is constituted by direct memory connections, such that our consciousness extends as far as we can remember ourselves. The revised psychological continuity theory consists of an overlapping strand of psychological connections, connecting successive stages of a self. Thus, it could be the

case that I am psychologically continuous with myself from a period of which I have no actual memories. Schechtman argues Locke's account is more conducive to a unified conscious subject; whereas the overlapping strand of memory connections is more likely to preserve things such as character and behavior, which are better understood as features of the moral self. This is because, as we saw in chapter 2 (section 3.3), Schechtman thinks that an account of a persons' persistence must provide an appropriate forensic unit to which further questions of moral responsibility (about the moral self) can be directed. According to Schechtman, this forensic unit corresponds with numerical identity, and should provide a determinate unit. She sees Locke's direct memory connections as more suitable for accommodating this concept of a forensic unit than the revised version of psychological continuity.

Another change Schechtman points out is the additional strands of psychological continuity—beliefs, desires, intentions, etc.—that are added to simple memory connections in a revised account of psychological continuity. Schechtman argues that, while it is clear in the case of memory connections why these should support the extension of the conscious subject, there is no good reason to incorporate these other psychological connections. Memories provide temporal anchors; they give us clear points of time to which we can trace the extent of our consciousness. Like the shift from memory connections to continuity, Schechtman sees these additional psychological connections as more suited to the continuity of personality or character than the sameness of consciousness. There are likely many beliefs, desires, and intentions that extend beyond the span of our consciousness, but if it is our consciousness that's important for a forensic unit, why should these other psychological connections matter? Also, these

additional psychological connections would be more prone to the kind of fading and splitting cases that challenge the determinate structure of numerical identity. Schechtman concludes that the only reason to incorporate these in an account of a person's persistence is if we wish to develop an account of both the forensic unit and the moral self, since she sees these additions as favouring the sameness of the moral self over the forensic unit.

These changes from Locke's memory theory to a revised account of psychological continuity, Schechtman argues, don't support the concept of a forensic unit that she sees as important for questions of a person's persistence. The switch from direct memory connections to an overlapping strand of memory connections (psychological continuity), combined with the additional psychological connections (beliefs, desires, intentions) don't support a determinate identity. Schechtman argues, at least Locke's direct memory connections are amenable to numerical identity in that they provide a unified conscious subject; revised accounts of psychological continuity lack this "unity". For Schechtman, this is a problem because she wants a determinate forensic unit which cannot be provided by contemporary accounts of psychological continuity.

Do these changes bring together the forensic unit and moral self, as Schechtman claims they do, and do they raise other concerns?

3.1.1 Does Schechtman Dismiss these Changes Too Quickly?

Schechtman's first point, for example, that psychological continuity, as opposed to psychological connections, is more in line with supporting the moral character than the forensic unit, seems misguided. As we noted in Chapter 1 (Sec. 2.2.1), the shift from

Locke's version of direct memory connection to the revised version of psychological continuity was directed specifically at Thomas Reid's Young Officer paradox and the problem with transitivity of identity. If you recall, the difference between psychological connectivity and psychological continuity is that connectivity consists in direct psychological connections, whereas continuity consists in overlapping strands of direct memory connections. It's worth keeping in mind that most cases of a person's persistence still entail direct psychological connections. The addition of psychological continuity, or an overlapping strand of psychological connections, is intended to support cases where the same person has persisted despite gaps in their memory chain—such as Reid's example. These cases are not intended to be the norm, and Parfit himself admits that both psychological connections and continuity are important: "We can agree that connectedness is not all that matters. Psychological continuity also matters. But we should reject the view that only continuity matters." (1984, 301-2) Parfit leaves an opening for both continuity and connectedness to matter, because in cases of fission, or fusion, we might want to place more importance on psychological connections. If we imagine a society where people are constantly fissioning apart and fusing together, then, somewhere down the line, everyone would be psychologically continuous with everyone else. In such a case, we would probably come to place greater importance on psychological connections as a way of distinguishing persons. Since these cases are important for Parfit's argument that personal identity doesn't matter in cases of survival, he argues that both psychological connectivity and continuity are important.

While I see why Schechtman wants psychological continuity to be limited to memory connections, instead of including the persistence of psychological

characteristics, I don't agree that an account of a person's persistence must provide a forensic unit as she understands it. I will get into the reasons for saying this more in chapter 4, but, essentially, I agree with Parfit that what matters in cases of a person's persistence is not numerical identity. Unless we're specifically trying to accommodate Schechtman's concept of a forensic unit as numerical identity, I see no reason not to expand the psychological criterion. For example, cases where memories have faded and are no longer reliable indicators of personal identity, but beliefs, desires and intentions persist. Also, patients with Alzheimer's: even though they might have poor memories, they may retain certain psychological traits that would support the notion that they are, in fact, the same person they were before their memories faded. Accepting Schechtman's interpretation of a forensic unit according to Locke's memory theory commits us to a very narrow understanding of a person's persistence that I see no reason to accept. If this view were true, then a person should be held responsible only for what they can remember. This leaves out all sorts of cases of repressed memory, aged memory, blackouts, etc. I'm not sure these cases are sufficient to claim that a person has ceased being the same person when they still possess psychological connections with a past stage of their self. I don't think there is much ground for reducing psychological continuity to memories.

We can see, once we look at these problems in more detail, there are valid reasons for why these changes were introduced to accounts of psychological continuity. Schechtman's claim that they favour an account that collapses the forensic unit and the moral self seems somewhat of a stretch. After these changes, there is still a forensic unit and a moral self. The only difference seems to be that this forensic unit doesn't

correspond with the concept of numerical identity in the way that Schechtman thinks it should. She takes these changes as contributing to what she sees as a decisive shortcoming of psychological continuity. Another change that contributes to this problem is that Parfit's account of psychological continuity is reductionist.

3.1.2 Reductionism

Being a reductionist account of a person's persistence is another feature of revised theories of psychological continuity that Schechtman sees as contributing to this problem. Being reductionist about a person's persistence means that a person's persistence consists in a series of facts. For an account of psychological continuity, these facts are simply psychological connections. There is nothing over and above these connections that constitutes a person's persistence. From this it follows that, if we are committed to psychological continuity, then questions about whether or not a person has persisted can be answered with reference to these facts. The further question of whether or not a person is identical with a past stage of their self becomes, ultimately, an empty question—a case where we can know all the possible facts relating to the matter without having a decisive answer. That is because there is no more information that is pertinent to this question than those facts that have already been accounted for by psychological connectivity/continuity. The further question of a person's identity becomes empty. This is because psychological connectivity is scalar (comes in degrees) and personal identity is determinate (all or nothing). There is no clear answer whether or not a person is identical with a past version of their self, instead the answer depends on where the parties involved in answering it determine the cut-off to be: at what degree does psychological

connectivity become sufficient/insufficient to support personal identity. This is a matter of convention, but there is no actual fact that supports the notion of personal identity.

According to Parfit, this indeterminacy is acceptable because identity is not what matters. What matters is Relation R, a case which he makes through the examples we looked at in section 1.4. For Schechtman, however, when we consider this reductionism in relation to the distinction between the forensic unit and the moral self, we run into problems. While Parfit's insight focuses specifically on what matters—i.e. the practical implications of psychological continuity; it also dispenses with the concept of numerical identity, which Schechtman considers to be necessary for a forensic unit and for answering the types of practical questions that concern a person's persistence. This is because, on a reductionist view, there is no unit of a person's persistence—just a collection of facts. For Schechtman, this rules out the concept of a forensic unit as numerical identity which she has in mind:

[for psychological continuity] there is no separate relation of identity which defines a more basic unity about which practical questions are raised. The identity of an individual person is made up of exactly the same relations we looked to to answer specific practical questions. The very idea of real unity is undermined by reductionism (2014, 34).

Thus, it would appear that if the forensic unit to which we ascribe a person's persistence is irreducible, and we are reductionists about a person's persistence, according to Schechtman, we find ourselves at an impasse. This is because she thinks both the forensic

unit and the moral self are necessary for capturing the sorts of practical concerns of a person's persistence, and a reductionist account eliminates the kind of forensic unit that Schechtman thinks is required. This problem, she argues, results in the extreme claim.

3.2 The Extreme Claim

We've mentioned the extreme claim already several times. Schechtman's main criticism of a reductive account of psychological continuity is that it results in what Parfit refers to as the extreme claim. This is basically the idea that, without a deeper fact of personal identity, our self-interest is irrational. Since this is a common criticism raised against accounts of psychological continuity and one that finds support whether or not you buy into Schechtman's analysis, it is worth discussing.

Parfit first refers to this worry as "the extreme claim" in *Reasons and Persons*, where he spends some time discussing its merits. Rather than deal directly with the claim that his account violates this principle of self-interest, Parfit opts instead to target the version of self-interest theory that supports the extreme claim. While we might not consider Parfit's responses exhaustive, he certainly provides the natural starting point for addressing this concern.

3.2.1 Parfit's "extreme claim"

Parfit introduces the extreme claim in Chapter 14, *Personal Identity and Rationality*, in *Reasons and Persons* (1984, 307). In this section, he discusses this prominent criticism raised against a Reductionist version of psychological continuity

from the perspective of self-interest theory. While there are various versions of this criticism, the main thrust of the claim is that a reductionist account of psychological continuity doesn't provide rational motivation for future directed self-concern⁴. The argument depends on a combination of Self-Interest Theory and a reduction of personal identity to a series of impersonal facts without any deeper connection.

On Parfit's account, Self-interest Theory is the view that the primary motivation for an individual is her own self-interest. "[F]or each person, there is one supremely rational ultimate aim: that things go as well as possible for himself. A rational agent should both have, and be ultimately governed by, a temporally neutral bias in his own favour" (Parfit 1984, 307). This is the theory that people are ultimately motivated by self-concern. When we combine this notion with a reductive account of personal identity, however, it becomes difficult to support. If a person simply consists of a series of facts—psychological or physical continuity—, then the further into the future we go, the less those connections hold, and the less reasons we would have for justifying any particular, future directed self-concern. The argument seems to be: since there is no deeper fact about personal identity, there should be no deeper concern for it either. If I learn someone is going to be in pain, there is some reason to prevent this pain; if I find out that it's *me* who is going to be in pain, it would seem I have additional reasons to prevent this pain. Since there is no deeper fact about personal identity, however, there is less reason to support these additional preventative measures. The further we set our sights into the future, the less significance the personal pronoun, *mine*, has with regard to denoting any particular concern for being in pain over and above the general concern raised by someone being in pain. So someone being in pain should hold no less significance for me

⁴ Parfit addresses variations raised by Butler, Sidgwick, Swinburne, Perry, and Wachsberg (1984, 307-8)

than the person with whom I will be psychologically continuous being in pain. Therefore, there is no particular self-motivated concern. There are several ways to cash out this argument.

Parfit distinguishes different variations of this concern depending on how much importance the proponent places on the additional facts of personal identity. Suppose that the persistence of a person does consist of some “deeper” fact. If this is so, then in a typical case of a person remaining the same person, we have biological continuity (the continued existence of the same organism), psychological continuity (the continued existence of the same stream of consciousness and so on) and personal identity (the deeper fact). If we are committed to the view that it is this deeper fact of personal identity that is important, then we would have to maintain that in cases of the continuity of a person, whether or not they are psychologically continuous holds no importance: personal identity with psychological continuity and personal identity without psychological continuity would be equally important. What’s really important is the deeper fact of personal identity, and not psychological continuity. This is the extreme version of the claim (hence the “extreme” claim). On this view, it is this deeper fact of personal identity alone that matters. The moderate version of this claim allows that some importance is placed on psychological continuity. Thus, if I left a psychological progeny to which I did not transfer the relation of personal identity, I still might have some special concern for its well-being. Consider Parfit's division scenario. In this case, a person is split into two psychologically continuous individuals. The moderate claim suggests it is possible to still have some future directed self-concern for the well-being of these psychological offspring, even if we shouldn't care for them in the same way as if they

contained the deeper fact of personal identity. Ultimately, Parfit thinks we should be able to settle on the moderate version of this claim, which offsets the worry that proponents of the extreme claim raise against reductive accounts of a person's persistence.

3.2.2 Parfit's Response

Early in his analysis of this concern, Parfit points out a distinction in how we might approach the Reductionist account of personal identity. "It is one question how, if we became reductionists, this would affect our attitudes, and emotions. It is another question whether, if the Reductionist View is true, these attitudes or emotions are justified." (Parfit 1984, 308) The distinction he's pointing out here is that, while a Reductionist view might entail a slightly different attitude towards the persistence of a person, this is not the same as suggesting it doesn't provide any rational support for our self-concerns. This difference can be accommodated by making the distinction between the extreme claim and the moderate claim.

Parfit admits of himself, that when he accepted the Reductionist view of personal identity, this drastically changed the way he saw himself, and his relationship with others.

When I believed that my existence was such a further fact, I seemed imprisoned in myself. My life seemed like a glass tunnel, through which I was moving faster every year, and at the end of which there was darkness. When I changed my view, the walls of my glass tunnel disappeared. Now I live in open air. There is still a difference between my life and the lives of other people. But the difference is less. Other people are closer. (Parfit 1984, 281)

If we accept the extreme claim, then we should have no special concern for our future self over and above anyone else. In part, this might not be such a bad thing. How much of what we consider self-interest would be better described as self-indulgence? Is there a deeper fact that supports this grandiose conception of the self, or is it perhaps more of a fiction than we'd like to admit? If so, then maybe the reality of self-interest isn't quite as justifiable as we have assumed. The problem with this view, however, is that there does seem to be good reason for, at least some, special concern for one's own future-self over others. This is where the moderate version of the claim comes in. It seems sensible to say that there is at least some special concern for a psychological progeny of yourself, even if she doesn't contain the deeper fact of personal identity. If, for example, you were to split into two psychologically continuous beings, neither would possess personal identity with you. If you are committed to the extreme claim, then you should have no special concern for either of these two individuals whatsoever. If, however, you accept the more moderate version, then, even though they might not possess that special concern of personal identity, there is still room for an extra measure of concern owing to the importance of psychological continuity. If we accept this version, then there is still room for self-interest on a reductionist account of psychological continuity, it just isn't as pronounced a concept of self-interest as we might have expected. But this fits with Parfit's position: if what matters in cases of personal identity is ordinary survival, and psychological continuity is just as good as ordinary survival, then that's it. There is no deeper fact to personal identity than psychological continuity, and we should have no more self-concern than we would have for a being that is psychologically continuous with us, because that's

all there is. The moderate claim supports future directed self-concern without having to resort to the deeper fact of personal identity that the extreme claim gets caught on.

Parfit favours this approach to the extreme claim, that once we introduce the moderate version, we have no reason to worry about a lack of future directed self-concern, since this view leaves room for it. He admits, however, that there is no decisive way to argue for the moderate claim over the extreme claim, so he proposes another solution, to discredit self-interest theory.

3.2.3 Self-interest Theory Argument

Instead of addressing the extreme claim head on, Parfit opts to argue against Self-interest Theory, effectively eliminating support for the extreme claim. His argument distinguishes the kind of self-interest that would reasonably result from his own view (that it is Relation-R and not personal identity that matter) with the kind of self-interest that would result from a non-reductive account of personal identity (a view where it is the deeper fact of personal identity that is important) to show why his own view provides the more sensible alternative.

Parfit's argument targets a central principle of self-interest theory. *The requirements of equal concern* is the notion that a rational person must be equally concerned about all parts of her future (Parfit 1984, 313). According to this principle, it is irrational for a person to care less about the distant future because it is more removed. A person might care less about a future event if the distance reduces the possibility of its occurrence, but she can't actually care less about her distant future self. On Parfit's

account, however, personal identity doesn't matter. What matters is relation R, which is psychological continuity and connectedness. Since psychological connectedness is a scalar quantity, it can come in degrees of importance. Accordingly, on this view, it is rationally justifiable to care less about one's further future since it is typically the case that psychological connectivity is weakened over long periods of time. Furthermore, it seems that people do, in fact, place more concern on their immediate future than the very distant future. If this is so, then the requirements of equal concern is not feasible, and neither is the version of self-interest theory that relies on it. What we think about self-interest theory will vary depending on whether or not we believe there is a further fact that supports personal identity. Understanding self-interest theory without this further fact seems to correspond better with the way we observe our self-interested behaviours. This supports Parfit's version over a concept of self-interest that calls for a strict concept of identity.

This is what Parfit has to say about the extreme claim. In the next section, we will look at how Schechtman's interpretation of the extreme claim comes out slightly differently.

3.3 Schechtman's Extreme Claim

The extreme claim is the main issue Schechtman sees as arising for contemporary accounts of psychological continuity. In her opinion, this occurs in reductionist accounts of psychological continuity because they rule out her specific concept of what a forensic unit should be in cases of a person's persistence. While the extreme claim as she presents

it is taken directly from Parfit, her criticism suggests she has a little more in mind regarding why it creates such a problem for psychological continuity.

One slight difference between Parfit's version of the extreme claim and the way Schechtman interprets the extreme claim has to do with the extent of the problems it creates for accounts of personal identity. In Parfit's explanation, the extreme claim is primarily focused on future directed self-concern: if there is no special fact about personal identity, then the further into the future a stream of psychological continuity goes, the less reason a person would have for caring about that particular future person over someone else. On Schechtman's analysis, however, there is no target for practical concern, self-directed or otherwise. Her interpretation of the extreme claim includes Parfit's version resulting from self-interest theory, but adds to it moral culpability: "the objection goes, [1] it does not make sense to hold people responsible for their past actions, and [2] there is no reason for a person to care about her own future well-being in a different way than she cares about the well-being of others" (Schechtman 2014, 35). So for Schechtman's version, according to a reductionist view of psychological continuity, not only would people not be motivated to pursue their own self-interests, but they couldn't be held accountable for their past actions, either.

For starters, Parfit's argument against self-interest theory wouldn't necessarily disarm Schechtman's version of the extreme claim that includes moral culpability. It doesn't seem to be the case that self-interest theory—at least as we have laid it out—is necessary in order to hold someone morally accountable for their past actions. A mentally handicapped individual who lacks a fully developed sense of self-interest and commits a violent action would likely be taken as responsible for that action: although we might not

hold them morally accountable in the same way as someone who is aware of the violence they are committing, they would still be restrained and have certain measures put in place to keep them from repeating this type of violent behaviour. In the same way, even if we were to assume the extreme claim, and say that a person has no reasons for prioritizing the welfare of their psychologically continuous progeny, there would still seem to be reasons for society holding particularly problematic, psychologically continuous beings as responsible for their actions. These seem to be two different sets of issues here: the prudential sense of self interest one might have for their future self, and the social instrumental concern that influences how people are held responsible.

The extreme claim, as Parfit lays it out, appears to be that there is no further fact of personal identity that supports the particular concern one might have for one's own future self over other future selves. This is different from the social concern that might arise from the actions of a being who is psychologically continuous. The first version (Parfit's version) seems to be that there is no difference between someone's future pain and *my* future pain. The second version (Schechtman's version), however, seems to assume there is no way of distinguishing who a person is, presently, from who they were in the past, and who they will be in the future. Just because there is no deeper fact denoting a deeper relation of personal identity, does not mean that we can't distinguish people at all. But this latter claim seems to be what Schechtman has in mind.

This reading of Schechtman's interpretation of the extreme claim is supported in various other parts of her account. "Relations that would justify the ascriptions of moral responsibility if they held within a forensic unit are not by themselves enough for such an ascription if the existence of such a unit is not presupposed." (Schechtman 2014, 35)

Schechtman places great emphasis on the necessity of a forensic unit to resolve psychological continuity's plight and avoid the extreme claim. One thing to ask here, is why she thinks that psychological continuity cannot provide a forensic unit in the first place? Even though this forensic unit might not conform with numerical identity in the way that Schechtman thinks it should, there is still a unit to be held morally accountable. Schechtman seems to think that the only way we can make sense of the practical concerns that relate to cases of a person's persistence is if that person provides numerical identity in the classical sense, in the form of a forensic unit.

According to Schechtman: "The extreme claim rejects the last part of [Parfit's] reasoning, denying that psychological continuity and connectedness as defined in these theories are important in the way that personal identity is important" (2014, 35). As we saw in Parfit's response, however, this interpretation generates the counterintuitive position that a person might be identical with a future person with whom they share no psychological, or biological continuity, whatsoever. Imagine a person A, splitting into two people B and C, all of A's psychological and biological continuity runs through to person B, but the deeper fact of personal identity goes on to person C. Now we have a case where A is identical with C, even though she shares psychological and biological continuity with B. This is a controversial outcome to say the least. This, however, can be avoided by accepting Parfit's moderate claim. On this version, even if we don't have the special concern of personal identity, we can still have a measure of concern for the future being with which we are psychologically continuous. Accordingly, there doesn't have to be a particular forensic unit that adheres to a strict concept of numerical identity in order for there to be a target for moral concern at all. For Schechtman, why she feels we need

numerical identity in order to have a forensic unit to which practical questions about the persistence of a person can be directed, is a question that she must answer if she expects her criticism to hold any weight. As far as I can see, however, we neither have, nor require, numerical identity in cases of practical concern.

Consider the following analogy: I am a Philosopher with an interest in personal identity, and you are a child with an interest in candy. I tell you that I'm interested in determining what it is that makes it such that a person is the same person over a period of time. And you tell me that you like lollipops. When I deliberate about what is of importance in cases of personal identity over time, I determine that it is, in fact, psychological continuity that really matters. You tell me that what you really like about lollipops is their taste. I decide to help you get to the bottom of this. So I ask if you like anything else the same way as you like lollipops. You tell me that yes, of course: you like gummy worms, skittles, and gobstoppers. Then I ask what if a lollipop came off of its stick, would you still like it? You say yes, it would be just like a regular hard candy; you'd just have to finish it in one sitting. I then ask what if the lollipop cracked in the packaging, and you had two pieces of a lollipop. You tell me that's fine, it happens sometimes, and then you simply have two candies, you can eat one now and save the other for later. It's like having two things you like instead of one—a double success. After this little discussion, I point out to you, that it seems what you really like is candy or flavoured fructose syrup. It doesn't seem to matter how you receive this candy, whether it comes in the form of a lollipop, skittles, or a gobstopper, provided it contains that sweet fruity flavouring, you're happy. You agree—this is what really matters. This candy doesn't have a single necessary form; it could take various forms, what's necessary

simply is that it be made of fructose syrup. And that's the same for the continuity of a person, what matters is that they are psychologically continuous, not that they take necessary the form of a numerical identity.

This analogy might seem glaringly obvious, but it's helpful for making sense of Schechtman's criticism. There is an important difference between having a necessary form, and having a form at all. Because candy can come in various ways, it doesn't have a necessary form. This doesn't mean that it doesn't have a form at all and that that form is not identifiable. You, the candy-loving child, can recognize various forms of candy: you mentioned lollipops, gummy worms, gobstoppers, etc. Even if I were to present you with something you were unfamiliar with, chances are you would be able to classify it as candy or not. Also, you can recognize certain candies as being the same candies, or different ones. You know, for example, the stick-less lollipop you've been sucking on is the same one you were sucking on several minutes ago, just as you know the broken piece of lollipop in a wrapper in your pocket is the same piece you put there this morning. You don't need to preface these candies with the possessive pronoun *mine*, or *yours*, to know whether or not they are the same. This can simply be distinguished relationally, without reference to any personal facts. It doesn't seem that much more needs to be said to derive any practical significance from your preference for candy.

Granted personal identity regarding future directed self-concern is somewhat more complex than tracking candy. However, psychological continuity provides a clear link between the person you were yesterday, the person you are today, and the person you will be tomorrow. Even if the person you were yesterday and the person you will be tomorrow are not identical with the person you are today, there is certainly good reason

for you to be responsible for the past self with which you are psychologically continuous and concerned about the future self with whom you will be psychologically continuous. I'm not sure how you could insist that only a relation of numerical identity could confirm this self-directed concern. Even if a future version of yourself has no direct psychological connections with your current self, that future self would still be connected with you in an intimate way, via psychological continuity. I'm inclined to believe that such a connection warrants at least some degree of special concern over and above the concern you might have for just anyone.

For Schechtman, the worry is that once one dispenses with the notion of numerical identity, we lose a target for moral concern. This, however, isn't the case, there is still a substantial form to psychological continuity, it just doesn't take a single, necessary form. If you recall our earlier discussion in chapter 1.1, we distinguished between numerical and qualitative identity. Objects can be numerically and qualitatively identical. If we dispense with the identity conditions, here, we are still left with quantitative and qualitative characteristics. Things cannot not have quantitative characteristics—otherwise the thing in question wouldn't exist. Just because psychological continuity is necessary for a person's persistence, doesn't mean a person wouldn't continue, only that continuing has a person doesn't take on a particular form. It simply means a person doesn't necessarily have a particular substantial presence. What's necessary is that a person possess psychological continuity; that psychological continuity can take on whatever form it happens to take on, be it one person, two people, a person with reduced psychological connectivity, and be tracked and held responsible according to that form.

3.4 Summary

The extreme claim is the worry that a Reductionist view of psychological continuity gives us little reason to justify future directed self-concern. There are several ways of dealing with this claim. Parfit points out that the distinction between a reductionist view of psychological continuity altering the way we treat personal identity and being justified in holding our beliefs about personal identity, are separate points. It's not unreasonable to say that if we accept a reductionist view of psychological continuity, our views might drastically change. We might come to see personal identity as not, in fact, mattering as much as we might have taken it to on a non-reductive view. This does not mean we have no reason to have concern for our future selves; this would simply mean our concerns might change. We should still be concerned about our future selves the same way we would be concerned for a being who shares our psychological continuity. What is it that we have in mind when we envision a being with whom we are psychologically continuous? Part of this is that we can conceive of a future person that will be psychologically connected with the person we currently are. We imagine a person who will remember and share our psychological connections with us in the same way that we do with the past stages of one's self. This is, certainly, an intimate connection, even if it doesn't contain a deeper fact of personal identity.

If this argument is still taken to be a problem, Parfit proposes undermining Self-interest theory by suggesting that a non-reductive account of personal identity, in fact, proposes the more unrealistic scenario that we should care about all future parts of ourselves equally. Psychological continuity, on the other hand, or relation R allows that

we care more for more closely psychologically connected versions of our self, and less for those that have fewer psychological connections, in the distant future. This seems to fit the way we think.

Schechtman's version of the extreme claim is slightly different. She insists that the worry covers not only our self-interest directed at future selves but the notion of being held accountable for past actions. This interpretation suggests several difficulties in the way that she interprets the extreme claim. The first point, that there is nothing to be held responsible, extends the notion of a person's self-interest to providing support for the basis on which a person should be held accountable for their actions. Even if without a deeper fact of personal identity, however, it would seem there are still grounds for holding a person responsible for their past actions according to psychological continuity. A deeper worry for Schechtman seems to be that when psychological continuity dispenses with numerical identity, it abandons a target for moral concern. Although this target might not take the necessary form of a forensic unit as numerical identity that Schechtman has in mind, there is still, nonetheless, a target. There is something that can be held accountable for its actions, and this is enough to meet the practical requirements that relate to a person's persistence. In no way is there a worry that psychological continuity doesn't have a substantial form. Once this is cleared up, the extreme claim ceases to pose a problem for reductionist accounts of psychological continuity. Since most of Schechtman's criticisms have been cleared up, we can now go on to look at her own account, the Personal Life View, which will be the focus of our next chapter.

Ch. 4 – Schechtman's Account

Schechtman's own account is in response to what she perceives as the need for a theory of personal identity that provides a suitable target for our practical concerns. To satisfy this need, she suggests that a theory should present a unified entity to which we can direct all of our practical concerns. These practical concerns should be inherently, and not accidentally connected with personal identity. Such a theory would constitute what Schechtman calls a dependence model. On her account, this model would reconcile a metaphysical account of personal identity with our practical concerns in a state of "mutual dependence between the two" (Schechtman 2014, 41). Accordingly, personal identity depends on practical concerns, since personal identity provides a suitable target to which our practical concerns can be directed; and our practical concerns depend upon personal identity, because personal identity must be present in order for practical judgments to be made. In order to provide the distinction between metaphysical and practical concerns as they relate to the concept of personal identity over time, Schechtman proposes what she terms literal identity

This chapter will present the Personal Life View (PLV) and look at some problems it faces. The first section will outline the view, explaining how it addresses the question of personal identity over time. It will also expand on the different elements of PLV, and Schechtman's use of Richard Boyd's Homeostatic Property Cluster (HPC) to tie them together. The second section will assess Schechtman's concept of literal identity, and how it provides an alternative to Parfit's version of psychological continuity. This comparison raises concerns about the suitability of Boyd's HPC for literal identity. The third, and final, section will look at more general issues arising from a broad construal of

personal identity and the HPC concept. Under scrutiny, it appears that these concepts lead the discussion back to some of the more traditional positions found in the debate.

4.1 Personal Life View

Schechtman's account of what constitutes personal identity over time is called the Personal Life View. She proposes this view, which will be expanded shortly, as a response to what she perceives as inadequacies in other accounts of personal identity. These accounts, she claims, fail to capture what is meant by personhood in a real, practical way. She argues that an account of personal identity should pertain to what we "literally" mean when we refer to persons—their literal identity.

4.1.2 Literal identity

Schechtman makes a point of defining her view as an account of literal identity. Although literal identity is similar to personal identity, she wants to distinguish it in several ways. She attributes the need for this distinction to the different commitments found in questions of personal identity. These questions can be divided into two types. The first type are basic questions about whether or not a person is the same person: is Person A the same as Person B? This type of question, she argues, has a definitive answer: either "yes" you are the same person I was speaking with yesterday, or "no" you are not. These are questions about the literal identity of a person. Since this adheres to the notion of numerical identity, it is strictly at odds with Parfit's claim that identity doesn't matter.

The second type of question corresponds to the different meanings we associate with the concept of personhood. This is similar to the practical concerns of personal identity discussed in chapter 2.3.1. When we refer to someone being a person, or the same person, we might have any number of concepts of personhood in mind: we might be thinking whether they have the same personality, are the same legal entity, or the same biological organism. These are all different ways to conceive of personhood, and there are different practical questions that correspond with these different conceptions.

Schechtman refers to these as questions about the moral self—although they are not all moral in nature, my guess is she generalizes the term “moral” to include any relational, non-intrinsic, concept of personhood. Unlike questions about literal identity, questions about the moral self do not always have definitive answers. Answers to these questions depend on a network of external relations and contextual interpretations.

This difference in concepts of personhood creates the problem of multiplicity outlined in chapter 2.3.2. This is the problem of bringing these different concepts of personhood and the moral self together within a single identity. “If we include all of the other practical relations that make up our form of life, the project of finding a single relation that defines the appropriate locus for our practical concerns appears, at least initially, hopeless” (Schechtman 2014, 81). The solution Schechtman proposes rests on the two-part division of judgments of accountability. The first part is ascription: where do we focus our inquiry? The second part is moral responsibility: is the focus of our inquiry morally culpable? For Schechtman it is the first part of this judgment, the ascription of moral concern, which is important for questions of literal identity. The second part, determining whether the target is morally accountable, concerns the moral self. This

second part takes into consideration all of the contingent factors that affect our decision; however, we cannot begin to consider these factors without first determining a suitable target for moral concern—this is the forensic unit we discussed in chapter 2.3.3. For Schechtman, the forensic unit is the key to establishing a connection between questions of literal identity and questions about the moral self. What we “literally” mean by personal identity—our literal identity—provides a suitable target for the kinds of practical concerns that deal with everyday questions about personal identity⁵. This, Schechtman argues, provides a necessary connection between the various concepts of personhood and literal identity: literal identity must provide an appropriate target for practical concerns relating to personhood.

Schechtman attaches several other stipulations to her concept of literal identity. Because literal identity must pertain to a forensic unit—a suitable target for moral concern—it must provide a unified locus for practical considerations. She reasons that this forensic “unit” must, inherently, be a singular entity (2014, 159). She also points out that literal identity would have to be irreducible—“the very idea of real unity is undermined by reductionism” (2014, 34). Since according to a reductionist view, as we saw with Parfit, there is no such thing as identity in the case of a person’s persistence, Schechtman automatically rules out the notion of reductionism. Also, because literal identity refers to numerical identity, it must be determinate, or all-or-nothing (2014, 32). Either person A is identical to person B, or she isn’t. Thus, an account of literal identity

⁵ While many of these concerns relate to personhood in a forensic (moral) capacity, Schechtman notes that there are other, practical concerns that do not fall under the concept of a forensic unit, which her account of literal identity is open to (2014, 76). While the forensic unit applies strictly to moral issues, she wants to expand this concept to include all of the ways that we conceive of personhood in practical speech—which is our “literal” identity. So the forensic unit and moral self are meant to explain the relation between literal identity and practical concerns, the latter of which expand the dimensions of the former beyond just moral concerns to include any and all practical concerns of personhood.

provides a suitable target for moral concern that is a singular, irreducible, determinate, unit. We can now move on to Schechtman's account of literal identity: the personal life view (PLV).

4.1.3 PLV

The Personal Life View (PLV) is Schechtman's account of literal identity. This view holds that a person exists for the duration of a single person-life; and that people are distinguished according to the specific person-lives they lead. A person-life consists of what constitutes a typical person's life. Obviously, there is a great variance amongst the different sorts of lives people lead, and determining what constitutes a "typical" person's life is contentious. Schechtman, however, focuses on the generic features that the more standard cases of person lives have in common.

According to PLV, a typical person's life follows a developmental trajectory. We begin as infants with limited physical and cognitive capacities, and occupy social roles that are, for the most part, passive. PLV differs from accounts that define a person according to their capabilities: for Schechtman, it is both the capacity to act as a person, and be acted upon as a person, that determines our status (2014, 69-71). Being treated as a person, even if one doesn't display many of the typical behaviours of a person, accords that individual social inclusion in what Schechtman refers to as person-space. Thus, infants, children, and individuals in persistent vegetative states (PVS) qualify as persons, even though they might lack certain capabilities that are associated with a fully developed adult.

On Schechtman's account, a fully developed person embodies all of the things we mean by personhood. A person should be sentient, self-conscious, a rational and moral agent, and a self-narrator (2014, 112). This combines biological features of sentience, and psychological notions of consciousness, with practical and moral aspects of rationality and agent-hood. To accommodate these features, Schechtman divides a standard person's life into three parts, though she notes these distinctions are artificial (2014, 112). The first part consists of the physical and psychological capacities of the individual. This part exists at the individual level and incorporates those features that are unique to the individual's physical and psychological makeup. The second part consists of the social behaviours and interactions that make up the individual's everyday life. These are things like personal relationships, projects, activities etc. This level takes into consideration the basic relations that influence a person's life. The third part is what Schechtman calls the social and cultural infrastructure of personhood. This includes the larger, social and cultural, structures and institutions that enable the development of personhood. Thus a person life is composed of these three interconnected parts: individual capacities, activities and interactions, and social infrastructures (Schechtman 2014, 115). These three parts can be understood as falling within three domains of personhood: the psychological, the biological, and the social. To accommodate these elements under literal identity, Schechtman uses Richard Boyd's concept of a Homeostatic Property Cluster.

4.1.4 Homeostatic Property Cluster

The Homeostatic Property Cluster (HPC) is a concept developed by Richard Boyd for the purpose of classifying species as natural kinds. This concept provides an

alternative to the traditional analytic model of defining natural kinds according to necessary and sufficient conditions. While the concept is developed specifically in relation to the biological classification of species, Boyd points out that he favours it being applied to many different areas of philosophy (1999, 145). The idea behind HPC kinds is that a natural kind can be defined by reference to a cluster of interrelated properties.

According to Boyd, HPCs occur in nature when properties dependently cluster together. These properties occur in a metaphorical (sometimes literal) homeostasis. It seems that, for Boyd, the term homeostasis is used quite loosely and simply refers to the tendency for certain properties to group together and form dependent relations. For example, when properties A and B occur together they might be conducive to property C; and, when properties A and C are found together they might be conducive to property D; so, it would not be uncommon to find properties A, B, C, and D all grouped together. Such clusters can be defined as a natural, or HPC, kind (1999, 143-4). Because of this, they don't have analytic definitions. Rather, they are defined naturally by the properties of which they are composed. This allows for imperfect homeostasis, where some properties might be absent. For example, you might find properties A, C, and D together, or properties A, B, and C grouped together and consider any 3 of the 4 properties enough to classify the HPC kind ABCD. The HPC concept also presumes that there will be extensional indeterminacy regarding these categories: cases where we might know all of the properties of an entity, but still be unable to determine whether or not it qualifies for membership within an HPC category. According to Boyd, settling cases of indeterminacy should be done a posteriori, and can generate controversial, theoretical debate.

Schechtman uses this cluster concept to account for how the different features of a person's life can be brought together under literal identity. A person's life involves the working together of biological, psychological and social functions. The way these different aspects relate to one another corresponds with the homeostatic mechanics of HPCs. This stresses the interrelation of these properties as a key feature of personal identity. These properties influence and encourage each other so that, in the case of persons, they occur together. Using this concept for defining person lives offers an alternative to dealing with borderline cases that have troubled traditional accounts of personal identity. Rather than looking for a single condition to determine whether a person has persisted, an HPC model looks at the relations between these conditions. So in cases of PVS, dementia, cloning etc., HPC models aren't confined to a single criterion for evaluating personhood. The HPC concept might seem well suited for developing a standard account of personhood, however, it runs into problems when this concept is applied to Schechtman's idea of literal identity.

4.3 Literal Identity and HPC

In our discussion of literal identity (4.1.2), we noted some of the key features that distinguish Schechtman's account of literal identity from other accounts of personal identity, particularly Parfit's version of psychological continuity.

Probably the most notable feature of Parfit's psychological continuity is his claim that in cases of personal identity, "*identity is not what matters*" (1985, 215). We discussed this claim in chapter 1.4. Basically, Parfit reasons that, if we accept an account of psychological continuity, then what's important for the continuation of a person is

psychological connections. Psychological connectivity, however, is a matter of degree. Numerical identity, on the other hand, is one-to-one. Thus, Parfit abandons the identity condition, concluding that what's important for survival is just a matter of psychological continuity, or Relation R. This also allows for branching cases where an individual might be psychologically continuous with two, or more, people.

Literal identity, however, is explicitly committed to numerical identity.

Schechtman makes this clear at the outset of her work: "Literal identity is like numerical identity in logical structure. It does not permit of branching or admit of degrees" (2014, 9). Her support for this seems to rest on the definition of a person as a unit of practical concern: "[t]his is because a person is, by definition, a unified locus of our person-related concerns, and there is no such unified locus where concerns are spread between two individuals" (Schechtman 2014, 159). For Schechtman, the social and cultural infrastructures of personhood are "built around the unity of the loci with which we interact", and wouldn't support personhood in any other way (2014, 164). To make this case with regard to Parfit's branching scenario, she points out that what results from cases of fission would not, in fact, capture what's important for the continuation of a person. Parfit's account focusses on the intrinsic concept of a person. Schechtman claims, however, that much of the importance of personhood is derived relationally. These relations accord specifically with a person as a unit of practical concern. Without such a unit, they would cease to pertain to personhood in a meaningful way. Would the relationship a child shares with her father endure having the original replaced with two psychologically continuous beings? Schechtman concludes these types of intimate personal relations would not endure fission in the sense of what's important for personal

identity (2014, 163). “If fundamental person-specific relational properties are taken into account the relation of the fission products to the original person is *not* the same as the relation of the unique survivor; fission is *not* the same thing, only doubled” (Schechtman 2014, 162). Thus, she is committed to the view that literal identity must provide a determinate, all-or-nothing relation. This, however, raises issues for the use of Boyd’s property cluster, and opens her account to several criticisms.

4.2.1 Determinacy and Clusters

When Schechtman makes the point that literal identity is similar to numerical identity in that it must provide a determinate, all-or-nothing relation, this raises concerns about the suitability of a HPC. HPCs specifically provide an alternative to analytic definitions, and so provide a less fixed concept of persistence. One of the defining features of the HPC concept is its extensional indeterminacy. This stands in contrast to the necessary determinacy of literal identity. While this isn’t exactly the same issue—identity (i.e., something determinate) can still occur within an extensionally indeterminate category—this relation needs explaining.

Because Schechtman’s account of literal identity is an account of identity, literal identity must be determinate. Schechtman considers identity to be fundamentally important for the sorts of practical concerns that relate to personhood. The ways in which we respond to the practical concerns of personhood (relating to the moral self) might themselves be indeterminate; however, questions about literal identity must have a clear answer (Schechtman 2014, 39-40). Given all the facts, we might still disagree whether someone should be held morally accountable, but there should be no question whether

s/he is an appropriate target for moral concern. Since Schechtman's account seeks to reconcile the practical concerns of personhood with the concept of literal identity, it must fulfill identity in this strict sense.

Boyd's HPC concept, on the other hand, is explicitly extensionally indeterminate. This is because it does not incorporate the necessary and/or sufficient conditions required to render extensional determinacy: "there will be many cases of extensional indeterminacy, which are not resolvable given all the relevant facts and all the true theories" (Boyd 1999, 144). Questions about the extension of a kind, he proposes, should be handled a posteriori, and often present a "difficult theoretical question" (1999, 143). This is a sensible way of going about the matter and accommodates the natural deviations that often do not conform to set categories. Here, then, there appears to be an incompatibility between Schechtman's account of literal identity, and Boyd's HPC cluster concept.

While there are cases where a category might be extensionally indeterminate, and a member of that category have a determinate identity, there would also be cases where an entity's membership would be indeterminate. On Boyd's account, extending membership to the more borderline cases would be a matter of a posteriori reasoning. If a conclusion is reached, then one might say that such an entity now has a determinate identity. Here, there is the possibility of having extensional indeterminacy of a category and a determinate identity of a member of that category, coextensively. However, it is also possible that if the matter were to remain unresolved—which HPCs allow for (Boyd 1999, 144)—this would not be a case of determinate identity. For Schechtman's view, this presents a real problem. In this regard, these two concepts—literal identity and HPC

concepts—are completely at odds. How Schechtman missed this, or why she didn't explain this relationship in more detail, is puzzling. On her account, we can't have something that's status of personhood is uncertain, nor should such a status waver depending on contextual interpretations. For literal identity, being unable to decide whether an entity meets the requirements of personhood (or continuation as the same person) would have severe implications: this is the central component of her theory upon which all of her practical concepts of personhood hinge. Thus, having extensional indeterminacy for the concept of personhood seems to present a real problem for literal identity. Another problem that comes up for Schechtman's account of literal identity has to do with which properties go into the cluster of personhood.

4.2.2 Homeostasis

One feature of the HPC concept that is underdeveloped is the nature of the homeostatic mechanism. This is because any specific definitions would exclude the concept's flexibility, which is why Boyd suggests the specifics would have to be determined a posteriori. This works for the more general concept of HPC; however, for PLV, the homeostasis of the concept of personhood would benefit from being developed. In Schechtman's defense, this is what the three-part division of a person-life, the psychological, biological, and social components of PLV are geared towards. These domains, however, are sufficiently different to warrant concerns about which types of entities might qualify for personhood when one is absent—which is entirely possible according to an HPC model.

It's worth noting that there are two separate questions regarding personhood according to PLV. This distinction is between what qualifies as a person, and what qualifies as a continuation of the same person. For Schechtman, since both questions are resolved in reference to a person's life, both are contained in cluster concepts. Because these concepts are understood according to a HPC model, they are equally flexible: there wouldn't be necessary and sufficient conditions in either case. Schechtman, however, makes a point of addressing these concerns separately.

To show what might constitute personhood according to PLV, Schechtman presents several examples of atypical cases found in the literature. Individuals with PVS and dementia, or children with cognitive developmental issues, for example, are considered persons according to PLV. As we saw in section 1.2, PLV accommodates these types of lives in person-space by focusing, not only on an individual's capacity to act as persons, but to be treated as one. Thus, while individuals suffering from certain conditions might lack the necessary cognitive or physical capacities to actively participate in personhood, they are still passive participants, which is sufficient for Schechtman.

The outcome of these cases according to PLV, however, raises other concerns. What about pampered pets that are treated much like persons? Or, cases of oppression where humans aren't treated like people? In both cases, Schechtman maintains that human qualities play a role in determining personhood. No matter how much we might treat our family dog like a person, we don't really think of it as a person: we don't expect it to speak, go to school, get a job, etc. (Schechtman 2014, 119-25). Though the same might be said of a cognitively impaired human, the difference is that, by default, we

expect these things, and only through learning the specifics of their condition, come to modify our expectations. For a dog, however, we do not hold these sorts of expectations. Similarly, with oppression, although the oppressor might not treat the oppressed as a person, they still recognize their personhood—according to Schechtman, this is why slave owners had sexual relations with their slaves, taught them religion, etc. (2014, 126-7). There are, however, exceptions to human personhood. She discusses non-human persons with respect to hypothetical talking, intelligent animals, Mr. Peabody the talking dog from *The Rocky and Bullwinkle Show*, for example (Schechtman 2014, 132). In such a case, PLV would deem Mr. Peabody a person, since he provides a suitable target for forensic concern. Here, we can see the flexibility that PLV allows.

Despite Schechtman's insistence to the contrary, however, what PLV qualifies as a person seems arbitrary. There are plenty of borderline cases that present difficulties for her view. Imagine, for example, a mannequin placed in a department store entrance, dressed to look like a person greeting customers. Obviously this is not a person, but it certainly could be accorded a place in person-space: people might wave to it, hold the door for it, move out of its way. An elderly man might tell a story about how rude the young gentleman was standing in the doorway that didn't help with his oversized parcel. Similarly, internet programs or telephone recordings designed to scam unwary victims: people believe these are people, they treat them as people, and, in some cases, these programs act and behave like people. The point isn't that, given all the facts, anyone would actually confuse them with real people, but that there doesn't seem to be a sufficient way of distinguishing these cases according to PLV. If the properties that constitute personhood are not necessary and/or sufficient, then what is the difference

between a severely impaired child who lacks the cognitive and physical capabilities of personhood, but is socially acknowledged as a person, and a computer program, that lacks the same cognitive capacities, but is also acknowledged in such a way? In certain cases, Schechtman falls back on the qualification of being human; however, she dispenses with this in the hypothetical case of Mr. Peabody, referring instead to the social interactions in which he is engaged.

In these cases, the influence of the biological, psychological, and social components in determining personhood is inconsistent. Must all three be present? Clearly not, since certain cases lack one aspect entirely—Schechtman, for example, discussed the possibility of androids or robots qualifying for personhood, which would have no biological component whatsoever (2014, 135-6). Does having a higher degree of one of these features compensate for a lack of others—it would seem so, since simply being a human is sufficient to qualify as a person, even without any psychological or social capabilities. Determining personhood in these cases seems ad hoc. Given these examples, without providing a clear logic, there doesn't appear to be a way for PLV to rule out some very counterintuitive cases of personhood.

These previous examples are related to the constitution of *personhood*. PLV's treatment of fission, or branching cases, is an example that relates to a person's *persistence over time*. Schechtman determines that fission cases would not constitute cases of the same person, according to PLV. Here, the question isn't whether the successive products would both be considered people; but whether they would be the same person, or contain what is important for the continuation of the same person, in any meaningful sense. This is because, for Schechtman, much of what is important in

maintaining a person's life requires there to be a unit of moral concern—the literal identity. The sorts of projects in which people are involved, and the institutions that they interact with, understand personhood as a unit. On Schechtman's account, this couldn't be otherwise. This, however, generates a contentious position.

Schechtman notices the worry of conventionalism that arises from this claim—the fear that personhood will amount to a socially constructed concept—and argues that the interrelation of the social, biological, and psychological features of personhood keeps this worry in check (2014, 119). Her position, however, still assumes that social infrastructures determine the qualifications of personhood. This, however, could just as easily be explained by individual developments at the biological and psychological levels determining the social institutions of personhood. For example, if fission cases were to occur, can we actually say social institutions would fail to consider these to be cases of the continuation of a person? More likely, social institutions would adapt to accommodate fission. Schechtman speaks to this in her analysis of fission cases (2014, 159), though it is still unclear why her view must be committed to the unity of personhood when it is so flexible in other respects.

The idea behind social infrastructures being inclined towards the unity of personhood raises other interesting considerations. Cases of collective agency, for example, occur where multiple individuals work towards a common goal and are legally treated as a unit. Or, corporations that have many of the legal rights, and in some ways, are treated like people. We might also consider family units, where minors exist under the same legal and financial status as their guardian; or countries, where the citizens are held accountable for the actions of their government. These are all cases where units of moral

concern comprise multiple persons. These units are not persons, but they are forensic units, and under such a protean account of personhood as Schechtman's, there is a legitimate question of why some of these cases wouldn't qualify for personhood.

We can see how having such a broad account of personhood raises concerns about the types of entities that might qualify for personhood. On the other hand, providing too clear of a criterion for literal identity doesn't fit with the notion of a property cluster.

4.2.3 Privileged properties

This last point brings up an important issue with the suitability of HPCs for developing an account of literal identity. HPCs specifically provide an alternative to traditional analytic definitions in terms of necessary and sufficient conditions. Accordingly, none of the properties are fixed, but rather the HPC accommodates fluctuations in the concept it defines. Schechtman, however, explicitly insists that literal identity must constitute a forensic unit – a locus for moral concern. For Schechtman, the property of being a forensic unit is a necessary feature of personhood. Due to the structure of an HPC, however, such a property, like any of the other properties, would not be fixed. An HPC account of literal identity would have to leave room for the possibility that things might change. This however presents a genuine problem for literal identity, which is explicitly committed to a fixed definition.

There are several ways Schechtman might respond to this claim. Boyd's account of HPC leaves an opening for the sorts of contextual, historical individuation of an object to remain consistent despite changes in its cluster of properties. Schechtman could argue that being identified as a forensic unit is part of the definitional integrity of personhood,

and serves a functional role. She might then go the route of claiming that anything that does not fulfill this role—of being a forensic unit—does not qualify for personhood. This, however, still raises the concern of why use an HPC to define a term that contains necessary properties.

One of the clauses of Boyd's definition of HPC concepts (with regards to natural kinds) speaks to this:

10. No refinement of usage that replaces t (where t is the identity of the kind) by a significantly less extensionally vague term will preserve the naturalness of the kind referred to. Any such refinements would require either we treat as important distinctions which are irrelevant to causal explanations or to induction, or that we ignore similarities that are important in just these ways. (Boyd 1999, 144)

This suggests that the natural kind will not be preserved if a fixed definition is attached to it. Which would be the case if Schechtman suggests that being a forensic unit should be understood as the functional designation of a person. It's not immediately clear how Schechtman intends to apply the HPC concept to her theory: whether she considers personhood to be a natural kind, or she simply wants to use HPC as an alternative to an analytic definition. If she does consider personhood to be a natural kind, then giving it the necessary quality of being a forensic unit directly contradicts this; and, if she simply wants to use HPC as an alternative to an analytic definition, but include necessary and/or sufficient conditions in that definition, it is counterintuitive. In either case, there seems to be a fundamental difficulty with using the HPC concept to define something that has

fixed properties. This feature of the HPC concept, that it is defined by reference to an indeterminate cluster of properties, raises one final point of contention regarding the irreducibility of literal identity.

4.2.4 Reducibility

As we noted in the outline of literal identity, one important stipulation that Schechtman includes is that literal identity is not reductive. The reason she includes this is that, in her view, any account of personal identity that can be reduced and defined with reference to impersonal facts, as with Parfit's account of psychological continuity, falls victim to the *extreme claim*: the argument that, without a distinct unit for practical concern, self-interest becomes irrational (Parfit 1984, 307-12). This is because, accordingly, we would not, in fact, be a self, but merely a collection of facts, or for Parfit, psychological connections.

Parfit, however, has good reason for insisting that any account of personal identity must be reducible. The concept of reducibility he is referring to is physical, as opposed to theoretical, reducibility. This is the idea that anything physical can be reduced into the particles of which it is composed, and these particles correspond with basic facts about the physical world. On this account of reducibility, anything physical can be reduced. Accordingly, if personal identity cannot be reduced, then it cannot be physical. Parfit argues this would, then, commit the proponent of irreducible personal identity to substance, or Cartesian, dualism. This is a sensible conclusion that many would consider uncontroversial, which, then, leaves Schechtman's account of literal identity committed to dualism.

If we move to look at Boyd's HPC concept, this clearly provides a reductive account for natural kinds. According to this model, HPC kinds are defined by the properties of which they are composed. This is precisely a matter of reducing a kind into its various properties. Even though the definitional identity might be distinct from the properties that make it up—so the HPC might maintain a fixed definition while its properties are in flux—this does not say that the entity itself is not reducible to those properties. It would seem that HPCs are a reduction, which presents a puzzle for how an HPC model might accommodate Schechtman's literal identity.

One possible suggestion is that the relation among the properties that constitute the HPC constitutes something over and above the properties themselves. This would give the HPC a formal element, and a structural unity that would allow one to say it is not merely reducible to the properties of which it is composed. This, however, discredits the concept of physical reducibility. Obviously, the facts that Parfit has in mind include processes, functions, etc.—otherwise, nothing would be truly reducible. If Schechtman means for her account of literal identity to be irreducible in this classic sense, which she makes a point of saying—“a definition of personal identity in terms of individual connections...is necessarily reductionist, and can provide a determinate identity only artificially...the question of who is the subject [literal identity]...require a more determinate kind of answer” (2014, 37-8)—then she must develop this idea more. If she means that her account should be irreducible in a different sense, then she must make clear what that sense is, and why this would motivate her critique of Parfit's reductionism. The problem is that Parfit's reductionism is necessary to avoid dualism with regards to the self. If Schechtman disagrees specifically with Parfit's reductionism,

then her view would entail dualism about the self. If, however, she has something different in mind, she needs to explain what that is in order to avoid charges of dualism. She would also have to explain how she accounts for a substantial concept of identity without committing herself to dualism.

This concludes our discussion of the problems with using an HPC concept to explain Schechtman's account of literal identity. In the next section we will present a more general look at the PLV concept of personhood, and HPC concepts.

4.3 Broad Views of Personhood and HPC Concepts

From the above analysis, we can see there are several points on which HPCs and literal identity are incompatible. This section will develop a more general discussion of Schechtman's construal of personhood, and HPCs as an alternative to analytic definitions. These remarks will conclude our discussion.

4.3.1 The Functionality of Analytic Definitions

There is certainly an appeal to providing an inclusive definition of personhood. It would seem that such a definition is better suited for the ways that we refer to persons in everyday speech than the more analytic definition of psychological continuity. Less specific concepts, such as Schechtman's PLV, are difficult to pin down and cover a broad range of possibilities. Such definitions also tend to be more versatile, which is why Boyd's concept of a property cluster has been such a popular way of considering natural kinds in the sciences. It would seem, however, that there are limitations to the usefulness

of these kinds of definition, as there would seem to be merits to definitions in the more analytic tradition.

In her work, Schechtman refers to cases where more specific definitions might be called for with regard to determining life and death in medical procedures (Schechtman 2014, 145). Boyd is quite clear about the functional and a posteriori research that would have to be done in order to classify natural kinds, and where applying a more specific definition would cease to pertain to a natural kind (1996, 143-4). Both accounts leave room for more specific definitions regarding functional processes. Specific, analytic definitions certainly are important for practical purposes. In a way, the difference between cluster concepts and analytic definitions proposes a trade-off: the more we dispense with technical analytic definitions, the more a posteriori research would have to be conducted, and events categorized according to their unique properties. This could, theoretically, continue until a point where categorization was dispensed with entirely, and everything would be assessed individually. Ultimately, this would be more precise, but much less practically efficient. The alternative is that we sacrifice precision for efficiency, and use more or less imperfect categories. This has proven to be a functional way of theorizing, and is certainly useful in the way we, as people, understand the world. This reasoning brings us back to some of the positions that first sparked this discussion.

4.3.2 What Goes Around, Comes Around

When we attempt to develop a practical account of personhood, we can see that there are certainly reasons for wanting a more analytic definition, featuring necessary and sufficient conditions. This is because it provides clear-cut answers on which we can base

functional procedures. Imagine, for example, an ER surgeon consulting a person's narrative history, and evaluating their place in person space in order to determine whether they should be resuscitated—this would be senseless. The difficulty with providing a more inclusive account for a concept such as personhood, is that when it comes to answering practical questions, the issue will likely have to be reduced to more analytic concepts, anyways. This is because such an inclusive account does little to differentiate the types of concepts that are applicable in real world scenarios.

Giving a broad definition of personhood simply shifts the meaning of personhood to more practical categories, which, essentially, makes the concept of personhood, itself, useless. If we are treating the biological, for example, then we address the concept of the human animal. If we are treating the rational, cognitive aspects of personhood, then we deal with the psychological. And if we are dealing with more practical, relational concepts of social interactions, we might look to the social. In this sense it is helpful to distinguish these different areas in accordance with our particular inquiry, and having concepts that specifically target these areas makes this division much clearer. Leaving these concepts as a cluster simply means more deliberation in particular cases. If, for example, personhood referred to all of the ways that we refer to persons in everyday speech, then my saying she ceased being a person would be meaningless until you ask me exactly what concept of personhood I have in mind. And, if you never bothered asking for clarification, this could be the cause of great confusion, if, for example, I had meant she became irrational and you took it that I that meant she was dead. The purpose of philosophy is to avoid this type of confusion and clarify these concepts, not further obfuscate them. Clearing up these concepts, however, brings us right back to some of the

more traditional concepts found in questions of personal identity: biological continuity, psychological continuity, and the moral self.

4.3.3 Summary

In this chapter, we have seen how some of the issues that Schechtman raised with regard to psychological continuity are not in fact as problematic as her account assumes. Particularly, reducibility and indeterminacy are well-supported features of psychological continuity. Avoiding these features and proposing an account of literal identity create many more problems than Schechtman acknowledges. Her best solution to avoid these problems would be to drop the concept of identity from her account entirely. We can also see how attempting to create a broad construction of personhood generates several inconsistencies, both in how she proposes to organize such an account, using Boyd's HPC, which seems problematic; and, in how effective such an account might be at actually addressing questions of personhood and persistence.

Incorporating the social domain of personhood into the question of personal identity is a very important feature of her account, and certainly affects the practical concerns of personhood. This relationship, however, is best understood in a causal, rather than constitutive relation. This distinction is an important point, which will be the focus of our next chapter.

Chapter 5 – Constitutive vs. Causal Personhood

In the last chapter, we saw how Schechtman's Person Life View provides a broad account of personhood by defining a person according to a person's life. A person's life, in turn, includes all of the activities and social infrastructures with which a typical person is involved. So for PLV, the identity of a person incorporates the external features that relate to personhood. As we noticed, it is difficult to develop the identity conditions for such an account of personal identity, and Schechtman's resort to a property cluster model presents problems of its own. The attempt, however, to incorporate relational and social components in an account of personhood is not completely mistaken. There does seem to be a legitimate case to be made for why we might consider the relations in which we are engaged to be a part of us. After all, humans are social creatures: we exist in groups, work together, and depend on each other in countless ways—in fact, the continuity of human life depends on these connections. You might find this reasoning compelling. Certainly, a great many people do. It is easy, I think, to be tempted by this view; but it's also easy to see how it can run into problems. Once we begin to admit these extrinsic relations as defining features of personhood, it becomes difficult to establish clear boundaries between the person and the causal systems in which they are engaged. This issue, however, can be clarified by a simple distinction between these different types of features of personhood: this is the distinction between the composition of persons and the causal role of persons. By making this distinction clear, we can see how our definition of personhood can accommodate the social domain of personhood without sacrificing a firm and functional definition of persons.

This chapter will focus on what we will refer to as the constitution of personhood (i.e., what makes up a person) and the causal role of personhood (i.e., the sorts of relations in which people are engaged). By making this distinction clear, I hope to shed some light on where Schechtman's account of personal identity has gone astray. First, I will begin with a look at the sorts of confusions that might arise from expanding an account of personhood to include its causal relations. To do so, I will look at an analogy Schechtman uses to support her position of developing identity conditions for a person life. I will then look at how this relates to the constitutive and causal accounts of identity. This distinction has played a significant role in the extended mind debate. In fact, the debate has turned on very much this point. I will examine this debate to help clarify how we should approach this issue with regard to personal identity. I will then address the third leg of Schechtman's triadic constitution of persons: the biological. I have already addressed this concern with regard to Olson, and so my closing remark will be brief. Finally, I will provide some concluding remarks on this essay. First, let's look at Schechtman's example.

5.1 The Storm Analogy

To make her case for how we might develop the identity conditions for a person's life, Schechtman borrows an analogy from Olson's biological continuity account. He argues that the conditions for determining the persistence of an organism amount to the continuation of the metabolic activities in which the organism's parts are involved. His point is that an organism doesn't have to be defined by having the same material parts, but by the processes, or events, in which those parts are engaged. Schechtman reproduces

Olson's storm analogy to make this point (1999, 139-40). Consider a tropical storm that hit Cuba last night and flooded Alabama today—to know whether this is the same storm or a different one doesn't involve any judgments about the material objects of which the storm is composed. We don't need to know whether in both cases the material parts are the same, to know that it is the same storm. Olson concludes that, if this is the case with storms, it should also be the case with biological lives—we needn't know all the parts, but simply the processes in which they are engaged, to identify a life. From this, Schechtman determines that (1) an entity can be identified by a life, and (2) that a life can be identified by a string of events (2014, 141). The important difference, however, is Olson's argument is strictly with regard to events that compose a biological life whereas Schechtman expands this argument to accommodate a person's life, which is a very different case.

Of the two claims Schechtman takes from this analogy, I don't find the first to be contentious: that an entity can be identified according to a life doesn't immediately seem problematic. The second claim, however, that a life can be identified according to a string of events is more contentious. For starters, it's quite vague about what sorts of events we're meant to consider and how they stand in relation to each other. For Olson's argument, these events refer specifically to the metabolic processes, or biological functions, in which the parts of an organism are involved. This gives a clear scope of the extension we can expect when delineating such an entity. For Schechtman, however, these lines are not so clear.

Remember that Schechtman's account of a person, PLV (chapter 4.1.3), incorporates many relational features of personhood falling under three different

categories: psychological, biological, and social. On her account, a person is partly comprised of the activities in which she is engaged and the social infrastructures that inform her daily life. It is one thing for Olson to define a biological organism according to the metabolic processes in which its parts are engaged. John's heart pumping blood to his kidneys, for example, is a particular event that would be highly unlikely to occur anywhere other than within John's body. It is a very different thing for Schechtman to define a person's life according to the relations and events in which they are engaged. Jess' relationship with her grandfather, for example, is an intimate relation that greatly influences her personal life. If we were to go about identifying Jess according to this relation, however, we would find it difficult. As it turns out Jess has a sister and several cousins, all of whom share similarly intimate relationships with their grandfather. If we are using this relationship as a defining feature of identity, then it becomes unclear whose identity it is that we have in mind: Jess, her sister, or one of their cousins? Another thing to consider with using this relationship to define identity is its impermanence: eventually Jess' grandfather will pass away; Jess, however, will go on living for much longer. It would be incorrect to continue referring to this relationship as a defining feature of Jess' identity, since it is a relationship in which she is no longer involved. Jess, however, is still very much the same person. Although she has certainly been affected by the death of her grandfather, she hasn't undergone any substantial changes that might give one cause to say she is no longer Jess. Using this relation as a defining feature, however, might lead one to say otherwise. This confusion results from identifying something according to the external events in which it is engaged, which is not a reliable process.

Consider the following scenario: a serial killer has gone on a killing spree. We identify this killer by the bodies s/he leaves, scattered across a relatively contained geographical area. There might be no biological indication of this identity, but there are certainly behavioural and social events that point towards it. After all, we know that a necessary feature of a serial killer is that s/he kills people, and these dead bodies would seem to indicate this behaviour. It would, however, be a mistake to treat this as a conclusive indication of identity. There might, in fact, be multiple killers acting within the same area. Similar reasoning can be applied to Olson's storm. It might be that the storm that hit Cuba yesterday is the same storm that flooded Alabama today, but there might also be two storms, or the storm might have split. Simply noting two or more proximate events is not sufficient for supporting identity. And basing our practical judgments on this reasoning could be problematic. If for example, I was a resident of Haiti that had just learned that the storm that hit Cuba yesterday was now flooding Alabama, I might relax and go for a walk on the beach. This would have tragic consequences, however, if the storm had split, and one of its offshoots was about to hit the very beach on which I planned to walk. Olson actually mentions this in his discussion of the analogy: "There will be hard cases of course, where it isn't clear whether we have two different storms or the same storm again...[t]hat is because storms (unlike lives) are not especially precise or well individuated events" (1999, 140). For Olson's account, this might apply, since he is referring specifically to events occurring within the clearly defined boundaries of a biological organism, but Schechtman doesn't have recourse to such "precise or well individuated events" since she proposes such a broad account of what constitutes a person life.

The main difference here is that Olson's account applies to internal events in an organism. These are events or processes that occur within an individual organism. Tracking these events can proceed over a spatial-temporal continuity. Schechtman's account, however, incorporates external, relational events. Without providing reference to a spatial-temporal progression, such as in Olson's organism, there is no guarantee that we are tracking the same thing. These are events that might occur between one or more individual organisms. So, we can see then why Olson's account lends itself to developing the identity conditions for a life much more readily than does Schechtman's. In Olson's case, there is a clear definitional boundary that can be tracked: that of the biological organism. For Schechtman, the boundaries of personhood incorporate biological, psychological, and social, attributes and activities, which aren't nearly as clear. I don't think that the external events in which persons are engaged, according to PLV, are always reliable indicators of identity. This distinction between the internal and external features of a person is very similar to the positions found in the extended mind debate, which we will discuss in the next section.

5.2 Extended Mind Debate

In a 1998 paper entitled "The Extended Mind", Andy Clark and David Chalmers⁶ proposed a radical version of externalism about cognition. Their suggestion was that the borders of cognition should be extended beyond "the boundaries of skin and skull" to

⁶ It is worth noting that Chalmers is not as committed to the extended mind thesis as Clark, a fact that is remarked in a footnote of the original paper (1998). He has since abandoned the view, and in the introduction to Clark's *Supersizing the Mind*, gives good reason for rejecting the thesis entirely: "Still, I think there is one potential principled place where the opponent of the extended mind can resist. This is an appeal to the dual boundaries of perception and action. It is natural to hold that perception is the interface where the world affects the mind, and action is the interface where the mind affects the world. If so, it is tempting to hold that what precedes perception, and what follows action is not truly mental". (Chalmers, 2009)

incorporate parts of the environment that actively contribute towards cognitive processes (1998, 7). The basic idea behind this claim is that cognitive processes should be seen to include many of the tools that facilitate cognitive computation: things like pencils, rulers, calculators, etc. They make a case for this claim through examining a series of cases involving cognition facilitated to different degrees by the use of external devices.

The best example they provide is the case of Inga and Otto (1998, 12-3). Inga wishes to visit an exhibition at the Museum of Modern Art, and she recalls that the museum is located on 53rd street. Her belief that the museum is located on 53rd street is not an *occurrent* belief; she obviously believed it before consulting her memory, it was something she had stored away. On the other side of town, Otto suffers from Alzheimer's and relies heavily on his environment to navigate his day-to-day life. To help retain useful information he keeps a notebook, which he consults whenever he needs to figure something out. His notebook functions much like memory. When Otto sees an advertisement for the exhibition at the Museum of Modern Art, he decides to go. He consults his notebook in order to determine that the museum is located on 53rd street. Similarly, as in the case of Inga, you might say that Otto held this belief before consulting his notebook—it didn't simply occur to him while scanning over the pages, it was something he had stored away. Clark and Chalmers maintain that these two cases are “entirely analogous: the notebook plays for Otto the same role that memory plays for Inga” (1998, 13). This example is meant to show how our cognitive process can be taken to include tools that exist outside of the brain.

Clark and Chalmers refer to these as cases of *coupled systems*. In these situations, the mind interacts with an external device in a dependent way. They argue that in such

cases, the device can actually be incorporated into the cognitive system: the device plays an important role in how the process is carried out, it is actively integrated into the system and modifies behaviour in the same sort of way as typical cognition (1998, 8). If one were to remove the device from the system, the function of the overall system would suffer much the same way as if you were to remove a part of the human brain. Why then should the cognitive system not extend to these devices? If we accept this reasoning then, according to what they refer to in as the *parity principle*, Otto's notebook would be an example of a cognitive process taking place outside of the brain (2010). Clark describes this as follows: "If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is (for that time) part of the cognitive process." (Clark 2010, 44). This is the basic idea behind their position.

We can draw parallels between the claims that support the coupled systems thesis of the extended mind, and the claims that would support the extension of personal identity into the social domain. Without certain social structures, for example, personhood would not be the same. Such social institutions—laws, governments, and culture—are integral to the way we view and understand the role of persons; and, if we were to remove any of these institutions, persons would cease to function in the ways that we have come to expect them to function. Since these systems are so integral to the role of personhood, why are they not part of the concept of personhood? After all, personhood exists within biological, psychological and social systems, all of which seem to influence how it functions. Why shouldn't we simply extend the concept of personhood, then, to incorporate these systems? Clark and Chalmers offer an even more radical suggestion in

the conclusion of their paper: “What, finally, of the self? Does the extended mind imply an extended self? It seems so” (1998, 18). While Schechtman's account doesn't go so far as an attempt to extend the concept of personhood into its environment, it isn't much of a leap to tack on technological and environmental to her list of psychological, biological, and social components of personhood.

We can see how, in some ways, the extended mind hypothesis mirrors an account of personal identity that seeks to incorporate many of the relational features of personhood within the actual concept. While these positions are not exactly the same thing, they are quite close and confront a very similar underlying question: where should the boundaries of such concepts be settled? In the next section, we will look at some fruitful responses to the extended mind hypothesis, which will help us addressing the relation in which personhood stands to society in questions of personal identity.

5.2.1 Setting up the Boundaries

In a 2001 article entitled “The bounds of cognition”, Fred Adams and Ken Aizawa defend the more common-sensical view that cognitive processes do, in fact, observe the brain boundaries. Their argument isn't that occurring within the brain is a necessary mark of cognitive processes, but rather that cognitive processes just happen to be delineated by the boundaries of the brain. To make their case, they set up what they consider to be the defining features of cognitive processes.

This is important because if they were to simply argue that cognitive processes occurred only in the brain, they would be begging the question. What they want is to show that (1) there are certain defining cognitive features, and (2) these features just

happen to occur within the brain. They aren't committed to the claim that these features can only occur within the brain, but they do want to rule out Clark and Chalmer's idea that cognitive processes occurring outside the brain are as commonplace as typing in a calculator. To show this, they propose what they consider to be necessary elements of cognition.

The first condition they determine to be essential for cognition is that it possess non-derivative, intrinsic content. Symbols, words, and numbers are all given meaning by virtue of some social, or conventional, assignment. Their content is derivative: it depends on the meaning that it has been ascribed by such conventions. In contrast, cognitive processes have non-derivative content (Adams & Aizawa 2001, 48). We don't require any referential apparatus to determine the content of our thoughts. Certainly, there are layers of referential sub-text that can be applied to our mental content; but the content itself is non-derivative—it has meaning by itself. While this isn't an altogether uncontroversial claim, Dennett (1987), for example, asserts that there is no such thing as non-derivative content, we needn't go so far as entirely non-derivative content in order to draw a meaningful distinction between the processes that take place within the mind, and those that occur externally. Simply the cognitive process of applying new meaning to existing strings of symbols, as would be the case for Dennett, should suffice to distinguish our cognitive processes from other functional processes that take place within the world. Accordingly, the content of our consciousness wouldn't have to be entirely non-derivative, but could be representational, incorporating both derivative and non-derivative content, which would be sufficient to distinguish it from other, non-cognitive processes.

The second condition for cognition is that cognitive processes are causally individuated. This is the idea that the cognitive can be considered like other natural domains. Adams and Aizawa reference the classic notion that science aims to carve nature at its joints (2001, 51). The idea they put forward is that science has always been aimed at determining the underlying causal processes of the world and categorizing things in accordance with similarities in their causal nature. As science develops, these natural joints become more apparent, and we are able to better categorize and understand the natural world. For Adams and Aizawa, although the science of cognition is limited, there are certainly features that distinguish cognitive systems from other systems that might resemble them, and devices that perform non-cognitive roles. They cite the example of a chess grand-master and a chess playing computer program to point out the difference (Adams & Aizawa 2001, 52). While the player and the computer might perform similar tasks, the actual cognitive processes that go on in the grand-master's mind are much different from those that occur within the computer. For starters, the computer will run through millions of scenarios, whereas the grandmaster will run through a much more limited number. Even though both the cognitive and computer system might be unrecognizable on some functional levels (i.e., playing chess), there is a clear difference in the processing that goes on behind the scene. Adams and Aizawa insist that there is a significant difference between what qualifies as cognitive, and what doesn't. They conclude that for something to qualify as cognitive, it must adhere to the above two conditions: it must have non-derivative content, and be causally individuated in the right way.

Adams and Aizawa accuse Clark and Chalmers of committing what they refer to as the *coupling-constitution fallacy* (2010, 67-8). This pattern, they claim, occurs when an object is coupled to a cognitive agent. The fallacy is then to consider that that object actually constitutes part of the agent's cognitive system. If, for example, you carry around your smartphone and constantly use it to check and record dates, then, according to Clark and Chalmers, there is a legitimate case to be made that the information you have stored qualifies as dispositional belief, and that smartphone as part of your cognitive system. If it is so integral to your daily planning that you consult it whenever making plans to the extent that you have become dependent on it to function in at normal capacity, they would argue, it behaves just like memory and constitutes a legitimate part of your cognitive process. Adams and Aizawa (2001), however, point out that this conflates the concepts of coupling with constitution. Borrowing their example, the kidney filters blood, and this processes is causally connected with the heart, which pumps blood throughout the circulatory system. We would not, however, say that the circulatory system filters our blood, this is a process that is causally allocated to the kidneys, not the entire system (Adams & Aizawa 2001, 56). If everything in the world that was coupled came to constitute the thing it was coupled with, we would quickly have a very short list of things in the world. What determines where we distinguish these things is according to their nature, which leads Adams and Aizawa to the two-part definition of cognition mentioned above.

This is just a brief survey of the extended mind debate, but there are clear similarities with the extension of the mind and the extension of personhood. In the next

section, we will look at how these responses might be helpful in understanding the constitution of personhood and its causal connection with other processes.

5.3 How this relates to P.L.V

There are certainly similarities between the extended mind thesis and the notion of extending personhood, but the two issues also have important differences. We can't directly apply the problems Adams and Aizawa point out with regards to extending a cognitive system, to the concept of personhood. For starters, personhood seems to be a different kind of concept than a cognitive system, although, in part, this might just be terminological. Maybe a better way of framing this comparison would be between persons and minds. We don't typically think of persons as systems, but rather parts within a system. In the same way, a cognitive system seems to imply multiple parts contributing towards cognition, while a mind implies a cognitive unit. Although there isn't a real difference between using the term "mind" and the term "cognitive system", the latter seems more susceptible to the coupling-constitution fallacy. Along these lines, Schechtman's account of a person life begins to resemble more of a system than what we might typically conceive of when we think of persons. On PLV, persons exist as a product of psychological, biological, and social spheres, all of which "constrain and support one another" (Schechtman 2014, 115). For Schechtman, the concept of personhood exists within this feedback loop of biological and psychological attributes and abilities and social institutions. So you might argue that her account does resemble more of a system of personhood, than simply a part within a system. If we see it like this,

then we can begin to see how Schechtman might also have slipped into a version of the coupling-constitution fallacy.

Another reason we can't directly apply the extended mind debate to the issue of personal identity is that the two arguments are positioned slightly differently. For Adams and Aizawa, the goal is to establish boundaries for cognition, and so they set out to define what they take to be the "mark of cognition": the necessary conditions for cognition (Adams and Aizawa 2001, 48). We can't directly apply this same reasoning to the case of a person's persistence, however, since in the personal identity debate, it is precisely the "mark" of personhood that is in contention. If we were to say outright that personhood is defined by psychological continuity, and we can delineate personhood wherever psychological continuity persists, this would simply be repeating our thesis. Adams and Aizawa argue that the mark of cognition happens to exclude extracranial cognition, not that being intracranial is necessary for cognition. We, however, want to say that the social domain is not constitutive of personhood. And if we simply define personhood to the exclusion of a social component, this does nothing to further our point.

Despite these differences, however, there are some useful points to be brought over from Adams and Aizawa's arguments. The coupling-constitution fallacy, for example, is quite relevant to our discussion.

5.3.1 Persons and the Coupling-Constitution Fallacy

The coupling-constitution fallacy seems particularly applicable to the case of personhood and social infrastructure. When we conceive of the notion of personhood, it is certainly connected with a concept of the social dimension. We would probably maintain, for example, that a person is causally involved with the social dimension—that they affect, and are affected by their social relations. We might also want to say that representations of these social relations might constitute the content of a person’s psychological connections—so their memories, beliefs, desires, etc., might reflect the social relations in which they are involved. This seems quite straightforward, and I don’t think anyone would debate it. It is a very different claim, however, to say that a person is actually constituted by the social dimension.

The social dimension is quite a broad category. “Social” is an abstraction that refers to many levels of relational and institutional organization. As a concept, the social domain isn’t really coupled with anything. Coupling, in a sense, refers to more particular relations. You might, for example, say a person is coupled with courts in a legal system; or, a person is coupled with a basketball in a game, or a person is coupled with their partner in a family. In a more abstract sense, you might even say a person is coupled with laws, in a legal system; or a person is coupled with rules in a game. In these cases, there are systems of causal relations that connect persons with objects, institutions, ideas, and other people. And these systems are all representative of some aspect of society. It is, however, unlikely that in any of these cases we would consider these systems to actually constitute a person. You wouldn’t, for example, say a person is constituted by a psychologically continuous human being *and* a courthouse; or, a person is constituted by

a psychologically continuous human being *and* a set of abstract rules; or, a person is constituted by a psychologically continuous human being *and* several other psychologically continuous human beings. But if we are to claim that persons are constituted by a social dimension, that is precisely what we're saying.

This conclusion seems very counterintuitive. But if we are to understand society to constitute something, then we have to look at exactly what we mean by society. We might even frame this as the further question of what constitutes society: what is society made up of? It would seem that society is made up of institutions, relations, and people—primarily, society is made up of people. Effectively, then, saying persons are constituted by society amounts to saying that people are made up of more people. But this is completely senseless. It leads to an illogical circle. If society is constituted by people, how can people be constituted by society? The only way to begin to even make sense of this is to suggest we are referring to the *concept* of personhood, and not persons, which is problematic for different reasons.

If we say that the concept of personhood is constituted by society, there are several things we might mean. We might mean (1) that the concept is literally constituted by society, i.e., society is what makes up the concept. We might mean (2) the concept of personhood is determined by society, i.e., personhood is a matter of social convention. Or, we might mean (3) society is a necessary cause of personhood, i.e., we wouldn't have people without society. These are three ways that we might, charitably, understand society as constitutive of personhood, but we can see that none of these offers a particularly favourable interpretation. The first option (1) that the concept of personhood is literally constituted by society is similar to the examples above. The only difference is

that instead of talking about an actual person being constituted by social relations, we are talking about the concept of personhood being constituted by social relations. So when we conceive of personhood, the idea is an amalgam of all the possible social relations in which people are engaged. But this just turns the concept of personhood into a much broader concept, which, effectively, begins to resemble the concept of society. The second option (2) that the concept of personhood is determined by society suggests that a person is whatever society deems a person to be. Schechtman leans away from this option when she suggests that her view does not, in fact, entail conventionalism about personhood, which is precisely what this interpretation suggests (2014, 119). Nonetheless, if this is what she means, then her view faces a real threat of conventionalism; the idea that personhood is a social construct and there is nothing that naturally demarcates a person. Either that, or this is trivial, and we accept the claim that personhood, along with every other concept, is influenced by the society that developed knowledge of it; and so, according to this view, everything is constituted by society, which really doesn't say anything. Finally, the third option (3) that society is a necessary cause of persons also becomes irrelevant. To claim that society is a necessary cause of personhood, doesn't mean that personhood is constituted by society. This would be the same as saying that because my parents are a necessary cause of me, I am constituted by my parents. Or, because water and oxygen are necessary conditions for human life, water and oxygen constitute human life. This third option leads back to the coupling constitution fallacy, where we confuse causal relations with being constitutive.

5.3.2 The Slippery Slope of Extending Personhood

There is a very strong case to be made for the idea that once we begin to admit external relations of a given entity as being constitutive of that entity, we are drawing very blurry lines. If we accept the claim that an entity is constituted by the relations in which it is engaged, then we are confronted with the problem of determining where those relations stop. If we include certain relations, then why not include the relations of those relations, and the relations of those relations, etc. This is the slippery slope of extending the defining characteristics of an entity to include external relations. There must be a meaningful way of distinguishing where those relations end. This is typically thought to be the border between intrinsic and extrinsic features of an entity. If we propose to extend that border, then we must provide a clear boundary.

Consider a simple example of two people playing basketball. John and Jess are playing basketball. I define John as being a psychologically continuous human being that is playing basketball. There are other parts involved in this game of basketball: there is a ball, there is a court, and there is Jess. If the relationship of playing basketball is constitutive of John, then it must also be constitutive of these other parts: the basketball, the court, and Jess. There doesn't seem to be a meaningful way for me to say that playing the game of basketball is any more a part of John than it is a part of, say, the basketball. If I did want to say there is an important way for me to distinguish John playing basketball from the basketball playing basketball, I would have to make the further step of defining John as a person (say, a psychologically continuous human being). To make this step, however, I must give an intrinsic definition of John: I must define John by what he is, not

by the activities in which he is engaged. As soon as I make this step, I admit that playing basketball is not actually constitutive of John. If I don't make this step, however, I must admit that John is constituted by the relation of playing basketball, as are the basketball, the court, and Jess. If the relationship of playing basketball constitutes all these parts, then I am not really talking about John anymore at all, I am talking about the game of basketball. Somehow, we have succeeded in transforming the definition of John into the definition of a basketball game. That is because playing basketball incorporates more than just John; it incorporates the basketball, the court, and Jess. To say that playing basketball constitutes John is incorrect; it would be more appropriate to say that John, along with the basketball, the court, and Jess, constitute playing basketball. Either that, or you wind up with some counterintuitive notion along the lines of John being constituted by the game of basketball, including a basketball, a court, and Jess. The same is true for persons and society.

The appropriate way to phrase this relation is that society is constituted by people. To say that people are constituted by society will generate the counterintuitive notion that a person is a society. This type of conceptual monism—that everything is one and connected—prompts the question: why stop there? Certainly, it is true that people depend on, and are influenced by the society which they are a part of. What about our environment? Certainly environmental factors, also, influence people: what we eat is determined by the nutrition that's available in our environment, our daily schedule is set by planetary rotation, everything, effectively, is determined by our environment. If we were to accept the claim that persons are constituted by society, then why shouldn't we accept the further claim that persons are constituted by their environments, or society is

constituted by its environment. This simply expands the definition. We exist in a world where everything is relationally implicated with everything else. If we take one step down the causal chain, why not take another, and another after that? Take technology as an example. Technology is a product of human civilization, as it progresses, it affects persons in many ways similar to social institutions (you might even consider technology a social institution). People live longer, behave differently, evolve differently all because of technology. Are people constituted by technology? Most people would say no. But again this is the same as arguing that people are constituted by society. Once we expand the definition of personhood, there is no longer a meaningful way to distinguish it from the causal chains in which it is implicated, and since we exist in a causally connected universe, everything gets subsumed into this ever-expanding relational definition. To avoid this, we must draw clear distinctions

5.3.3 Carving Nature at its Joints

The point Adams and Aizawa (2001) make, that science aims to carve nature at its joints, is solid council for avoiding the confusion that might arise from misinterpreting definitional boundaries. When we attempt to identify the underlying causal processes of reality, definitional boundaries come to represent our causal understating. As our understanding deepens, these boundaries change: they might become more specific, or they might isolate different causal networks. This is how the natural world presents itself to us. Even if we say that persons and society are not really natural causal categories, there still appears to be something fundamental about these concepts that denotes a meaningful distinction.

When we talk about people, and when we talk about society, we mean two very different things. When we think about a person, most people probably think of a human animal; they might go so far as to distinguish a fully developed adult human animal. This would be a typical response. When we push this response further, we might refine the response to psychological continuity as being what defines a person. This is very different from when we think about society. When we think about society, most people probably think of a group of people. Is it possible for animals to exist in societies? Would we consider apes, for example, or ants, to exist in a kind of society? In part, this answer depends on how we define society. What about people: is it possible to consider a person existing without society? In either case, there seems to be a clear distinction we have in mind when discussing these concepts. And this gives us an approximation of where we might set up causal boundaries.

Even if we want to consider a person to be a necessary part of society, as Schechtman does: “[b]eing brought into the form of life of personhood may be described as being accorded a place in person space” where “person-space [is] the social and cultural infrastructure within which persons interact and which supports personhood” (2014, 114, 115). This doesn’t mean we can’t distinguish where a person ends, and society begins. If we couldn’t make this distinction, then we might as well do away with these categories entirely. A rough definition might be that society is composed of people, a weaker definition might be that society is composed of a group of animals living together (or something to that effect). If we accept the stronger definition, that society is composed of people, then we might be tempted to define a person [as a constituent of society. But this is circular: a person is a member of society, society is a group of people,

people are members of society, etc. We can see how this sort of explanation is problematic. It is true that people exist within the social causal system. The same, however, can be said of many other things. Things can exist within the same causal system without being constituted by that system. A more appropriate way would be to say society is constituted by people, and people are constituted by psychologically continuous human beings, and psychologically continuous human beings are constituted by memory connections and living organic bodies, etc. When we develop explanations, we usually try to avoid giving the explanans in the same terms as the explanandum, because this is circular. It's not incorrect; it's just a tautology.

An unfortunate consequence of this system of explanation is that it might generate the misbelief that causal processes follow a single direction just like explanations. We might infer from this that one thing causes another and there is no reciprocity. This would be a mistake. And the kind of mistake that might lead one to suggest that persons and society must, then, constitute each other. This mistake would arise from the belief that if people constitute society, there is no causal relation between society and people: people cause society and society doesn't cause people. If this were the case, then we could see why someone might mistakenly argue that society constitutes people. If we understand that there is a causal reciprocity, however, there should be no reason to be concerned with setting things out the way I have. Society simply is a group of people, and the causal direction goes around. I'm not sure anyone would disagree with this notion. It seems quite straightforward. Things affect the causal system they exist within and that causal system, in turn, affects them. This is the nature of a system. Once we distinguish the parts

from the system, we get a much clearer understanding of the causal and constitutive nature of things.

5.4 The Biological Component

If you recall from chapter 4.1.3, Schechtman's PLV consists of psychological, social, and biological features. So far I have argued in support of a view of psychological continuity that corresponds with Parfit's version. I have just finished outlining the types of difficulties that arise from an account that attempts to confuse the causal relations in which persons are engaged with the constitutive features that make up personhood. This deals with two of the three legs Schechtman proposes for P.L.V. Evaluating the biological component was not a major focus of this essay. I do think there is much to be said in this regard and that there are very intricate relations that exist between the concepts of persons and humans. Doing these relations justice requires a much more extensive survey than this essay offers. Nonetheless, I do think we have provided sufficient material to dismiss the third part of Schechtman's account. In this respect, I think some of the main distinctions between persons and human animals that we looked at in our discussion of Olson (chapter 2.2.1) hold up. Beyond that, there is a similar argument as the one we just made regarding the causal and constitutive distinction in personhood to be made regarding why the human animal doesn't necessarily constitute the person. By parity of reasoning, we can apply some of our arguments against the inclusion of society in the constitution of persons to the inclusion of biological organisms in the constitution of persons. I admit that this is a quick and dirty solution, and actually

flushing this out would be a much more complex endeavour, but it gives something to work on another time.

5.5 Concluding Remarks

Over the course of this essay, I have argued for a version of psychological continuity similar to the view put forward by Parfit (1984). Parfit stands up well to scrutiny, including Schechtman's challenges, Schechtman, on the other hand, does not stand up well to challenges to her point of view. Therefore, Parfit's approach is to be preferred. His is the view that a person's persistence can be accounted for by psychological continuity. When we think of ourselves as persons, precisely what we are thinking of that distinguishes us from merely being human animals is our consciousness; and, what makes it so that we can know that we are the same person we were at an earlier time is following that stream of consciousness back through those moments of which we are composed. This paints a simplified picture, in reality there is a complexity and depth of psychological connections that span one's existence as a person. Memory just gives us a clear sense of the trajectory through which we have passed. This psychological connectivity is what is essential to being a person.

Certainly, as persons, there are other important relations in which we are involved. These relations determine how we understand ourselves, and how we interact with each other and within the world in which we are situated. I have argued against biological and social components as being constitutive of personhood because I don't think these are essential. That is not the same as saying that I don't think these components are an

important part of the way we are. I think that much of the way we are and how we understand ourselves depends on precisely these relations. But I think that, before we can even begin to be a certain way, or understand ourselves as such, there must be something upon which that understanding is imprinted. That something is our consciousness, it is our psychological continuity. That is what is necessary for a person to exist, and continue existing over a period of time as the very same person she was when she first thought of herself.

Bibliography

- Adams, F., & Aizawa, K. (2001). The Bounds of Cognition. *Philosophical Psychology*, 14(1), 43-64.
- Adams, F., & Aizawa, K. (2010). Defending the Bounds of Cognition. In R. Menary (ed.), *The Extended Mind*. (pp. 67-80). Cambridge, Massachusetts: The MIT Press.
- Baker, L. R. (2000). *Persons and Bodies: A Constitution View*. New York; Cambridge, U.K.;: Cambridge University Press.
- Boyd, R. (1999). Homeostasis, Species and Higher Taxa. In R.A. Wilson (ed.), *Species: New Interdisciplinary Essays* (pp.141-85). Cambridge, MA: MIT Press.
- Brook, A. (2014). Tracking a Person Over Time is Tracking What?. *Topics in Cognitive Science*, 6(4), 585-598.
- Butler, J. (2008). Of Personal Identity. In J. Perry (ed.), *Personal Identity* (2nd ed.). (pp. 99-105). Berkley, CA: University of California Press. (Original work published 1736)
- Chalmers, D. (2008). Foreword. In A. Clark, *Supersizing the Mind: Embodiment, Action, and Cognitive Extension* (pp. ix-xxiv). New York: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis*, 58(1), 7-19
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. New York: Oxford University Press.
- Clark, A. (2010). Memento's Revenge: The Extended Mind, Extended. In R. Menary

- (ed.), *The Extended Mind*. (pp. 67-80). Cambridge, Massachusetts: The MIT Press.
- Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, Mass: MIT Press.
- Erikson, E. H. (1959). *Identity and the Life Cycle: Selected Papers*. New York: International Universities Press
- Erikson, E. H. (1968). *Identity, Youth, and Crisis* (1st ed.). New York: W. W. Norton
- Locke, J. (1741). *An Essay Concerning Human Understanding: In Four Books*. (The twelfth ed.). London: Printed for John Beecroft, at the Bible and Crown, over-against the New Church in Lombard-Street.
- Olson, E.T. (1997). Was I Ever a Fetus? *Philosophy and Phenomenological Research*, 57(1), 95-110.
- Olson, E. T. (1999). *The Human Animal: Personal Identity Without Psychology* (New ed.). New York: Oxford University Press.
- Olson, E. T. (2007). *What Are We?: A Study in Personal Ontology*. New York: Oxford University Press.
- Parfit, D.A. (1971). Personal Identity. *The Philosophical Review*, 80(1), 3-27.
- Parfit, D.A. (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Perry, J. (1976). The Importance of Being Identical. In A.O. Rorty (ed.), *The Identities of Persons*. (pp. 67-90). Berkley, CA: University of California Press.
- Reid, T. (2008). Of Mr. Locke's Account of our Personal Identity. In J. Perry (ed.), *Personal Identity* (2nd ed.). (pp. 113-118). Berkley, CA: University of California Press. (Original work published 1785)
- Schechtman, M. (1996). *The Constitution of Selves*. Ithaca, NY: Cornell University Press.

- Schechtman, M. (2014). *Staying Alive: Personal Identity, Practical Concerns, and the Unity of a Life*. Oxford: Clarendon Press.
- Shoemaker, D. (2009). *Personal Identity and Ethics: A Brief Introduction*. Peterborough, Ont: Broadview Press.
- Shoemaker, S. (1970). Persons and Their Pasts. *American Philosophical Quarterly*, 7(4), 269-285
- Shoemaker, S. (2008). Personal Identity and Memory. In J. Perry (ed.), *Personal Identity* (2nd ed.). (pp. 119-134). Berkley, CA: University of California Press. (Original work published 1959)
- Strawson, G. (2015). 'The Secrets of All Hearts': Locke on Personal Identity. *Royal Institute of Philosophy Supplement*, 76, 111.
- Williams, B. (2008). The Self and the Future. In J. Perry (ed.), *Personal Identity* (2nd ed.). (pp. 179-198). Berkley, CA: University of California Press. (Original work published 1970)