Fragments of Mind

A Thesis on The Subject of Experience,[1] *the Structure of Consciousness, and Split-Brain Patients*

A thesis submitted in partial completement of the requirements for award of the degree of B.Phil in Philosophy

Faculty of Philosophy

University of Oxford, 2023

Candidate Number: 1057910

Word Count: 24586

[1] The first part of the thesis title ("The Subject of Experience") reflects an earlier plan which didn't come to fruition. There's currently comparatively little discussion of the subject of experience. Due to administrative reasons, the thesis title couldn't be changed, as it was after the meeting of the GSC in week 5.

Thesis Abstract: Neuropsychological research on so-called *split-brain patients*—individuals whose corpus callosum has been severed as a treatment for severe epilepsy—have revealed the possibility of startling breakdowns in the transmission of sensory information between the cerebral hemispheres. These experimental results have often been interpreted as demonstrating that split-brain patients have two streams of consciousness. In opposition to such views, this thesis argues for a conception of split-brain patients as having a single stream of consciousness that's disunified in highly specific experimental settings but unified outside of those settings. It is also argued that conceptions of split-brain patients as having two unified streams of consciousness and conceptions according to which they have one consistently unified stream of consciousness rely on a mistaken assumption: That a set of experiences belong to a stream of consciousness if and only if all the contemporaneous experiences in the stream are unified with each other.

PROLOGUE 4
CHAPTER 1: PHILOSOPHICAL FOUNDATIONS 8
1 Consciousness and Streams 8
1.1 Consciousness 8
1.2 Streams of Consciousness 10
1.3 The Vehicles of Experience and the Contents of Experience 12
2 Unity and Cardinality 16
3 Multiplicity Principles 18
3.1 A Vehicular Multiplicity Principle 19
3.2 A Thought-Consistency Multiplicity Principle 21
3.3 A Experiential Consistency Multiplicity Principle 24
4 Forms of Unity 28
4.1 Spatial Unity 28
4.2 Temporal Unity 30
4.3 Phenomenal Unity 32
4.4 Access Unity 36
4.5 The Relation between the Spatiotemporal Unity and Phenomenal Unity 37
CHAPTER 2: FOUNDATIONS ABOUT SPLIT-BRAIN PATIENTS 40
5 Splitting the Brain 40
6 Experiments on Split-Brain Patients 43
6.1 How Split-Brain Experiments Work 43
6.2 Two Split-Brain Experiments 47
6.3 Split-Brain Patients and the Hemispheres' Capacity for Experience 50
6.4 Split-Brain Patients and the Malleable Stream View 51
CHAPTER 3: BEYOND MULTIPLICITY AND UNIFIED SINGULARITY 53
7 Split-Brain Consciousness and Multiplicity Principles 53
8 Split-Brain Consciousness and Sufficiency Singularity Principles 62
8.1 The Case against Unity 62
8.2 Bayne's Defence of Phenomenal Unity 72
CHAPTER 4: MULTIPLICITY AFTER ALL? 78
9 Multiplicity Through Lack of Unity? 78
9.1 Unity NSP and Multiplicity 78
9.2 Schechter's Argument for Multiplicity 79
9.3 Nontransitivity as a Counterexample to Unity NSP 83
10 Must Unity be Transitive? 84

10.1 Unimaginability and Inconceivability 85
10.2 The Weak Inconceivability of Nontransitivity 88
10.3 Weak Inconceivability and the Unity of Imaginative Awareness 91
10.4 The Uniqueness of What-It's-Likeness 93
EPILOGUE 96
BIBLIOGRAPHY 98

PROLOGUE

Looking back at a long and illustrious career as a neuropsychologist, Michael Gazzaniga describes the following as his greatest moment:

> Nothing can possibly replace a singular memory of mine: that of the moment when I discovered that case W.J. could no longer verbally describe (from his left hemisphere) stimuli presented to his freshly disconnected right hemisphere. An experiment I had designed, executed and carried out alone as a mere graduate student […] had worked. With it, the modern split-brain story was born […].

Gazzaniga's sense of pride isn't misplaced. For this experiment, along with others conducted by Roger Sperry and Joseph Bogen, helped sparked a revolution in research on the human brain. The experiment was about so-called *split-brain patients* (henceforth "SB-patients"); individuals for whom the corpus callosum, the fibre tract connecting the cerebral hemispheres, has been surgically removed, standardly as a way of treating severe epilepsy. What Gazzaniga's experiment showed was that, under certain carefully controlled experimental conditions, the two surgically disconnected hemispheres were incapable of interhemispherically transmitting information, and that they exhibited a significant degree of independence in the processing of sensory information and the execution of cognitive-behavioural tasks. A salient example of this, which Gazzaniga mentions in the quote above, is the left hemisphere's (LH) incapacity to describe the sensory stimuli presented to the right hemisphere (RH). Much of what today is known about the hemispheres' specialization for different kinds of tasks derive, either directly or indirectly, from experiments conducted on SB-patients.[2]

[2] An example is LH's specialization for language processing and the right RH's specialization for spatial reasoning (Gazzaniga & Sperry, 1967; Sperry, 1974).

Philosophers have been interested in research on SB-patients ever since the first emergence of the phenomenon. Research involving SB-patients may help shed light on traditional philosophical problems about the nature of personhood, agency, and consciousness. This thesis is concerned with the bearing of SB-patients on one such traditional topic, namely *the unity of consciousness*. It's common in the philosophical literature to hold that SB-patients put pressure on the notion that consciousness always is unified (Lockwood, 1989; Marks, 1980; Nagel, 1971; Puccetti, 1981; Tye, 2003). This thesis will attempt to answer the question of what the *structure* of split-brain consciousness is.

When asking this question, there are two conceptually distinct things one could have in mind. Firstly, there's a *counting* question: Do SB-patients have a single stream of consciousness? Secondly, there's a *unity* question: Are the experiences of SB-patients unified? I'll argue that the answers to these questions diverge. SB-patients do have a single stream of consciousness, whose experiences are unified given normal circumstances. However, under specific experimental conditions, SB-patients' experiences become disunified. This will be called the *Malleable Stream View*. On this view, SB-patients have a single stream of consciousness that is malleable with regard to whether all its elements are unified.[3]

For the purpose of this thesis, there are two main competitors to the Malleable Stream View. The first is the *Unified One-Stream View*. This is the view that SB-patients have a single stream of consciousness whose elements *always* are unified, even in the

[3] Philosophers who have defended something along the likes of what I call the Malleable Stream View are (Lockwood, 1989; Marks, 1980).

context of neuropsychological experiments (see, e.g., Alter, 2010; Bayne, 2008, 2010). The second major competitor to the Malleable Stream View is the *Two-Stream View* of SB-patients. This theory posits that SB-patients have two unified streams of consciousness. Versions of the Two-Stream View has been popular among neuropsychologists (see, e.g., Sperry, 1968, 1984), and also has philosophical defenders (Davis, 1997; Puccetti, 1981; Schechter, 2018).[4] I'll focus especially on Tim Bayne (Bayne, 2010) and Elizabeth Schechter (2018) as representatives of the Unified One-Stream View and the Two-Stream View.

I'll argue that the Unified One-Stream View and the Two-Stream View make a common mistake which the Malleable Stream View avoids. This mistake is assuming that a set of contemporaneous experiences can belong to a single stream of consciousness *only if all members of the set are unified with each other*. It will be shown that there aren't satisfactory grounds to accept that assumption.

This thesis is composed of four main chapters. The first chapter presents the fundamental concepts and principles which will serve to inform the discussion in subsequent chapters. In addition to elucidating the concepts of consciousness and the notion of conscious unity, the chapter describes sufficient conditions for experiences to belong to *different* streams of consciousness (*Multiplicity Principles*), sufficient conditions for experiences to belong to the *same* stream of consciousness (*Sufficiency Singularity Principles*), and necessary conditions for experiences to belong to the same stream of consciousness (*Necessity Singularity Principles*).

[4] Schechter is special in that she believes split-brain patients have two streams of consciousness and that they are composed of multiple subjects of experience, each with its own stream of consciousness (Schechter, 2018: ch. 2).

The second chapter focuses on neuropsychological research on SB-patients (henceforth "SB-experiments"). Particular attention will be paid to discussing experiments that philosophers have thought are relevant to the unity of consciousness.

In the third chapter of the thesis, I'll argue that that the experiences of SB-patients don't satisfy the antecedents of Multiplicity Principles, nor the antecedents of Sufficiency Singularity Principles. In connection with arguing for the latter thesis, I'll show that experiments on SB-patients establish that such patients have disunified experiences. Hence, the Unified One-Stream View is false.

In the fourth and final chapter of the thesis, I'll make the case that the disunities between SB-patients' experiences don't establish that the Two-Stream View is true. In other words, even though SB-patients' experiences are disunified, it doesn't follow that they belong to different streams of consciousness. The truth of the Two-Stream View is established by the fact of disunity only if the unity relation is *transitive*. I'll argue that there aren't any convincing reasons why the unity relation cannot be nontransitive.

CHAPTER 1: PHILOSOPHICAL FOUNDATIONS

1 Consciousness and Streams

This thesis is concerned with streams of consciousness. Two questions about streams will be answered in this section. Firstly, what kinds of mental states belong to streams of consciousness? And secondly, what are streams of consciousness themselves?

1.1 Consciousness

Stative and Episodic Mental States

A distinction can be made between stative mental states and episodic mental states. A *stative* mental state can exist over an extended period of time, and needn't be understood as having a definite occurrence or manifestation. Examples could be beliefs, desires, intentions, and other "standing" or "dispositional" mental phenomena. In contrast, *episodic* mental states essentially have the character of occurrences that exist for a specific duration of time. Examples could be sensory experiences, bodily sensations, onsets of emotion, episodes of imagination or thinking, explicitly entertaining a hypothesis, and so on. In short, while stative mental states *are*, episodic mental states *happen*.

Stative mental states admit of a broadly Freudian distinction between consciousness and unconsciousness. Just as there can be conscious beliefs or conscious desires, it's conceivable that there could be unconscious beliefs or unconscious desires. However, the distinction between consciousness and unconsciousness doesn't appear to be

applicable to episodic mental states. For example, sensations or particular episodes of thought aren't the right kind of entities to exist in a subject's unconsciousness.

This observation will serve as a blueprint for how the notion of consciousness will be understood in this thesis. Henceforth, the notion of a conscious state, or experience,[5] will be understood as the notion of an episodic mental state that can figure in a subject's stream of consciousness. Whatever matches this characterization is an experience.

Later in the thesis, the discussion will turn to two special forms of consciousness. The first of these is *phenomenal consciousness*. A mental state is phenomenally conscious iff there's something it's like to be in that state (Block, 1995; Chalmers, 1996). A phenomenally conscious mental state has a distinctively subjective character (cf. Nagel, 1974). For example, in having a phenomenally conscious experience of pain, one doesn't merely have a state that's caused by tissue damage, that results in avoidance behaviour, and which bears certain characteristic relations to other mental states (e.g., fear and anger). One also has a state with a subjective character, a character that's "[…] essentially connected with a single point of view […]." (Nagel, 1974: 437.)

The second special form of consciousness that will be of interest is *access consciousness*. A mental state is access conscious iff it's "[…] (1) inferentially promiscuous, that is, poised for use as a premise in reasoning, (2) poised for rational control of action, and (3) poised for rational control of speech." (Block 1995: 231, in-text citations omitted.) Thus, access conscious mental states are available for control by a subject's cognitive system, facilitating rational control of action, speech, and inference.

[5] "Consciousness" and "experience" will be treated as synonymous terms, as will "conscious state" and "experience."

1.2 Streams of Consciousness

Experiences aren't free-floating entities. Rather, they're organized together in certain systematic ways. It will be assumed in this thesis that experiences are organized together in *streams of consciousness.*

Streams of consciousness are processes of experiencing. Paradigmatically, they'll be composed of sequences of (sets of) experiences occurring in temporal succession. For example, at the present moment, my stream of consciousness is constituted by a succession of auditory experiences of furious typing and visual experiences of alphabetical characters appearing on a computer screen.

Three points must be made about how the notion of streams of consciousness will be understood in this thesis. Firstly, it will be assumed that any stream of consciousness has a subject. This means that for any stream of consciousness E, there's a subject of experience S who undergoes all the experiences $e_1, \ldots, e_n$ that feature as elements in E. There are various ways in which this assumption could be justified. One potential justification would be the Strawson-inspired view that streams of consciousness are individuated by, and owe their nature to, their subject of experience.[6] But the assumption could be defensible even given a broadly Parfitean (Parfit, 1984) conception of subjects according to which subjects are constructions out of agglomerations of experiences. Consistently with that conception, it could be true that if there's a stream of

[6] P.F. Strawson writes that "[states], or experiences, one might say, *owe* their identity as particulars to the identity of the person whose states or experiences they are." (Strawson, 1959: 97.)

consciousness, then there will be some subject-construction or other to which every experience in the stream can be attributed.

Secondly, the question should be asked what the relationship is between experiences on the one hand and streams of consciousness on the other hand. It was mentioned above that, paradigmatically, streams of consciousness are constituted by experiences. But can a single experience constitute, and thereby be identical with, a stream of consciousness?

The answer that should be given to this question depends on whether experiences have a point-like or an interval-like temporal structure. If an experience has a point-like temporal structure, then it occurs at a particular, point-like instance in time. On this view, experiences don't really have the character of processes or events that *unfold*. In contrast, if experiences are understood as having an interval-like temporal structure, then they occur not so much at a particular point in time as during a particular stretch of time. Here, experiences *do* have the character of processes, the character of being temporally unfolding events.

If particular experiences occur as point-like instances in time, then a single experience cannot, on its own, constitute a stream of consciousness (since streams of consciousness must be processes). Conversely, if experiences have an interval-nature, then there's no reason to deny that a single experience can constitute a stream of consciousness. On the other hand, even on the interval conception, there will be streams of consciousness that aren't identical with single experiences. For streams of consciousness that are constituted by multiple succeeding sets of experiences aren't identical with any single experience.

Thirdly, how long-lived are streams of consciousness? Specifically, can they persist through interruptions in experience, such as during periods of unconsciousness or dreamless sleep? It will be assumed in this thesis that they *can*. For example, if one loses consciousness as a result of falling off one's bike, and subsequently awakes, that doesn't involve the obliteration of one's old stream of consciousness and the emergence of a new stream of consciousness. Rather, in this case, one's stream of consciousness recontinues when one regains consciousness.

If streams of consciousness are entities capable of survival over periods of interruption, then, during those periods at least, streams of consciousness depend for their existence on other entities. It cannot be a brute, unexplained fact that a recontinued stream of consciousness is identical with an earlier, interrupted stream. One potential explanation of the possibility of persistence appeals to the subject of experience. According to this proposal, a recontinued stream of consciousness is identical with a previously interrupted stream of consciousness in virtue of the fact that the streams have the same subject of experience. However, there may be other explanations of the possibility of stream persistence.

1.3 The Vehicles of Experience and the Contents of Experience

Attention must be paid to the distinction between *vehicles* and *contents* of experience.[7] To elucidate this distinction, we'll initially examine the role of the vehicle/content distinction in the domain of public representations.

[7] Hurley is a typical example of a philosopher insisting on the significance of this distinction (Hurley, 1994, 1998, 2003).

Firstly, let's consider a postcard adorned by a photograph of the Eiffel Tower. It's natural to make a distinction between the object of depiction and the depiction itself. The Eiffel Tower is made of iron and was constructed between 1887 and 1889. In contrast, the postcard is made of paper and was printed a week ago in Croatia. Secondly, consider a sign in a fast-food restaurant, exhorting customers to "Be cautious; the floor is slippery when wet." Just like with the postcard case, one can distinguish between content and vehicle. The content of the sign is an exhortation to avoid reckless walking. In turn, the vehicle is a piece of plastic inscribed with alphabetical characters. In both cases, the object of depiction, or that which is stated, have different properties than the depiction, or the medium of the statement.

That isn't to say that the contents and vehicles of public representations cannot share any properties. For example, in the case of a photograph of a red apple, the object of depiction and the depiction itself shares a property, namely redness. But the two can still be distinguished.

Similarly to how it's possible to distinguish between the vehicle and content of public representations, it's possible to distinguish between the vehicles and contents of consciousness. The *vehicle* of an experience is an event in spatiotemporal-cum-causal reality. If one is a physicalist, one could identify the vehicles of experience with neurophysiological events in the brain. However, non-physicalists will presumably also want to distinguish between, on the one hand, *a subject's instances of experiencing*, and, on the other hand, *that which* the subject experiences, the contents of her consciousness.

Using a spatial metaphor, consciousness can be thought of as a bucket, or container. Anything that enters into the bucket qualifies as a content of consciousness,

while anything located outside the bucket doesn't so qualify. There are at least two ways of further elucidating what the contents of consciousness are.

Option 1: The content of an experience *e* can be identified with the objects, qualities, and relations that the subject *S* is experientially presented with when having *e*. For example, if *e* is a visual experience of a red tomato, then the content of *e* will be the features of the tomato that are visually presented to *S*; its redness, it's approximate sphericality, and so on. It's possible that some the objects, qualities, and relations in question are identical with objects, qualities, and relations in the mind-independent, external world. However, it's also conceivable that there are cases for which this isn't true, such as hallucinatory experiences.

Option 2: The content of an experience *e* can be understood as *e*'s propositional content. Interpreted in one way, Option 2 may be different from Option 1. For if by "propositional content," we mean "proposition that is contained in consciousness," then it isn't obvious that the propositional content of an experience *e* will include the objects, qualities, and relations that a subject is experientially presented with when having *e*. Rather, the propositional content of *e* would perhaps be better conceived of as a proposition that is about these objects, qualities, and relations, but which doesn't contain them as parts.

However, there's way of disambiguating Option 2 so that it has the same implication as Option 1 about what contents of experience are. For "propositional content" can also mean "that which a proposition contains." On a Russellian view of propositions, the constituents of a proposition are the objects, properties, and relations that are "contained" in it. Hence, the constituents of propositions about a subject's

experiences will be the objects, qualities, and relations that the subject is presented with when having the experiences. Let's call this way of understanding Option 2 "Option 2.2," in contradistinction to the view of Option 2 presented in the last paragraph, "Option 2.1." Option 1 and Option 2.2 give identical verdicts on what contents of experience are.

Therefore, Option 1, Option 2.1, and Option 2.2 are all consistent with understanding contents of experience as propositional contents. (Option 1 because its treatment of contents of experience is identical with that of Option 2.2.)

To summarize what has been said so far, the vehicle of an experience is an event in the spatiotemporal world, while the content of an experience are the objects, qualities, and relations a subject is presented with by having that experience, or the experience's propositional content. Thus, at least *prima facie*, there appear to be differences between the properties of experiences' vehicles and the properties of experiences' contents. For example, assuming that some form of physicalism is true, the vehicles of experience are spatially located inside the human brain. However, it's false that the contents of experience are spatially located inside the human brain. Conversely, contents of experience can be described as having certain qualitative features; for example, the redness of a ripe tomato. It's false that a vehicle of an experience whose content is a red tomato itself must be red.

These observations aren't presented as purported demonstrations that the vehicle of an experience must be distinct from its content. Rather, they're intended to show that it's intelligible to distinguish between contents of experience and vehicles of experience. The only assumption that will be made about the relation between contents of

experience and vehicles of experience is that it's conceivable for the two notions to come apart.

2 Unity and Cardinality

The discussion of the unity of consciousness in this thesis will be informed by two basic notions: *Unity relations* and *cardinality principles*.

Unity Relations

A *unity relation* is a binary relation between experiences.[8] Unity is best understood as a determinable relation between experiences that subsumes more specific determinate unity relations.

Cardinality Principles

Cardinality Principles are principles that can be used to determine the number of streams of consciousness in a given situation. They concern the number of streams of consciousness, hence their "cardinal" character. Cardinality Principles come in two main forms.

Firstly, there are *Multiplicity Principles*. A Multiplicity Principle specifies a sufficient condition for two experiences e_1 and e_2 to belong to different streams of consciousness. More stringently put:

[8] What's called "unity" here is often also called "co-consciousness."

General Multiplicity Principle: Necessarily, for any two experiences e_1 and e_2, if e_1 and e_2 satisfy condition C at t, then e_1 belongs to stream of consciousness E_1 at t, and e_2 belongs to stream of consciousness E_2 at t, where $E_1 \neq E_2$.

Secondly, there are *Singularity Principles*. These come in two forms. The first, Sufficiency Singularity Principles (SSPs), articulate *sufficient* conditions for a set of experiences to belong to a single stream of consciousness. The second, Necessity Singularity Principles (NSPs), articulate *necessary* conditions for co-membership of a stream of consciousness.

It's extremely compelling to suppose that if two experiences are unified, then they belong to the same stream of consciousness. If that's true, then there's a Sufficiency Singularity Principle with the following form:

General Unity Sufficiency Singularity Principle (General Unity SSP): Necessarily, for any set of experiences $e_1, \ldots, e_n$, if $e_1, \ldots, e_n$ are all unified in respect X at t,[9] then there's a single stream of consciousness E at t that contains $e_1, \ldots, e_n$ as members.

The specific form of this principle depends on what determinate of the determinable unity relation is substituted for "unified in respect X." There may be SSPs that don't

[9] The claim that "$e_1, \ldots, e_n$ all are unified in respect X" is the claim that, for any experience in the set, that experience is unified with every (other) experience in the set.

have anything to do with the unity relation. However, in this thesis, the only SSPs that will be discussed are Unity SSPs.

Likewise, the only NSPs that will be discussed are Unity NSPs:

> **General Unity Necessity Singularity Principle (General Unity NSP):** Necessarily, for any set of experiences $e_1, \ldots, e_n$, if there's a single stream of consciousness E that contains $e_1, \ldots, e_n$ as members at t, then $e_1, \ldots, e_n$ are all unified in respect X at t.

Just as with Unity SSPs, the specific form of Unity NSPs depends on what specific determinate of the unity relation is substituted for "unified in respect X."

3 Multiplicity Principles

In this section, three Multiplicity Principles are presented and discussed. The first two are taken from Peacocke (1983) and Williams (1978). They'll be discussed primarily to elucidate the difference between Multiplicity Principles about vehicles of experience and Multiplicity Principles about contents of experience. The third Multiplicity Principle, which articulates a non-contradictoriness constraint for contemporaneous elements of a stream of consciousness, is the Multiplicity Principle that will be discussed in connection with SB-patients later in the thesis.

3.1 A Vehicular Multiplicity Principle

In a section of *Sense and Content* (1983: 176-179), Christopher Peacocke presents a thought experiment involving a hypothetical SB-patient. Examining Peacocke's discussion of the thought experiment will reveal that he espouses a Multiplicity Principle which is about the *vehicles* of experience rather than the *contents* of experience. Peacocke presents the thought experiment as follows:

> Consider someone who has, temporarily, a split-brain because his *corpus callosum* has been severed. We can make sense of the possibility—and indeed it sometimes happens—that there are two distinct but qualitatively identical experiences grounded in (different) states of the same brain [---]. There is a clear motivation for saying that in the split brain there are two token experiences, a motivation which does not appeal to the identity of persons. The two experiences have different causes—for example, one results from the stimulation of one nostril, the other from stimulation of the other nostril—and they may have different effects too. (Peacocke, 1983: 176-177.)

Peacocke writes that it's conceivable that a SB-patient *S* has two simultaneous and qualitatively identical experience tokens e_1 and e_2. By the experiences' "being qualitatively identical," Peacocke means that they involve the same "what-it's-like" qualities. Put in the terminology that has been employed so far, e_1 and e_2 have the same content. The reason for positing two experiences here rather than one is because there are two vehicles of the pertinent experience-type.

Peacocke then goes on to argue that e_1 and e_2 must belong to different streams of consciousness. He does this by invoking the following principle:

> Token experiences, which are *total* in the sense that their type fully specifies what it is like (visually, aurally, etc.) for someone who has that experience, individuate minds or centres of consciousness. Distinct total token experiences at a given time must belong to different persons or minds; and there are no further distinctions between minds at a given time than are determined, directly or indirectly, by this principle. (Peacocke, 1983: 178.)

What Peacocke claims here is that, for any moment in time, there can only be one experience-token of a given kind in a stream of consciousness.[10] More specifically, at any moment in time, there can just be one *total* experience-token of a given kind in a stream of consciousness. A total experience-token is an experience of a maximally determinate type, a type that uniquely captures the qualitative character of the experience. So Peacocke's principle is consistent with the possibility that there can be simultaneous experience-tokens e_1 and e_2 of the same *determinable* type, provided that e_1 and e_2 differ in terms of their *determinate* type. An example of this would be if e_1 were an experience of scarlet and e_2 an experience of vermillion. What the principle rules out is that a single stream of consciousness can accommodate two simultaneous experience-tokens of the same maximally determinate type. Applied to the thought experiment, the principle entails that *S* has two streams of consciousness, streams that are "individuated" by e_1 and e_2.

We can formulate the principle Peacocke is relying on as follows:

> **Token Uniqueness Multiplicity Principle (TUMP)**: Necessarily, for any two experiences e_1 and e_2, if (i) e_1 and e_2 are simultaneous, and (ii) e_1 and e_2 are tokens of the same maximally specific qualitative type, then e_1 belongs to stream of consciousness E_1, and e_2 belongs to stream of consciousness E_2, where $E_1 \neq E_2$.

[10] Peacocke writes about "minds" and "centre of consciousness." That should cover "streams of consciousness" as well.

TUMP is a Multiplicity Principle about the vehicles of experience. For at the level of the contents of experience, there's no room for distinguishing between e_1 and e_2, as they have the same content.

While Peacocke isn't very explicit about why he subscribes to TUMP, the best explanation has to do with counting streams of consciousness. TUMP articulates a clear condition for counting streams of consciousness: There are at least as many streams of consciousness in a given situation as there are simultaneous experience-tokens of the same maximally determinate type.

3.2 A Thought-Consistency Multiplicity Principle

In his cogito argument, Descartes famously argued from the proposition "I think" to the proposition "I exist." In a response to Descartes, the 18th century German physicist Georg Lichtenberg argued that Descartes wasn't justified in positing a thinking "I"; at most, what Descartes had established, according to Lichtenberg, was that "there is thinking" ("cogitatur"). If Lichtenberg is right, then the first premise of the cogito argument isn't justified. That is, Descartes hasn't adequately defended his claim that there's a thinking subject.

In his book on Descartes, Bernard Williams (1978: 79-86) raises a problem for the Lichtenbergian project of constructing impersonal reformulations of putatively first-personal statements about thinking. Among other things, the problem has to do with contradictory thought contents and entailment relations between thoughts. What Williams means by "thought" here is perhaps better captured by the notion of

judgement or belief. Here, we're dealing with "thought" in the sense of taking a content or proposition to reflect actual states of affairs in the world.

Let's consider the following three statements:[11]

(1) It's thought: P.

(2) It's thought: ¬P.

(3) It's thought: P and ¬P.

Clearly, it's possible for there to be a thought that P, and a thought that ¬P, without there being a thought that P and ¬P. The only way to ensure this is by *relativizing* (1) and (2) to their own respective "thought-world." (Williams, 1978: 81.) If (1) and (2) don't occur in the same thought-world, it isn't guaranteed that there will be a thought world in which (3) is true. Thus, (1) and (2) could be true without (3) being true.

Naturally, the big question here is what "thought-worlds" are supposed to be. Subjects or persons are good candidates. However, this option isn't available for a Lichtenbergian who denies that there are subjects of experience. Williams considers various impersonal relativization strategies, and argue that they fail (Williams, 1978: 80-86).

Williams' argument assumes that it's possible for (1) and (2) to be true without (3)'s being true. It's on this basis he introduces the requirement that (1) and (2) should be relativized to different thought-worlds. However, there's an alternative way in which

[11] The exposition has been adapted, and doesn't entirely follow Williams.

Williams could have made a similar point. Here, the starting point is (3) rather than (1) and (2). Is (3) a possible thought?

Some thoughts with contradictory contents are may be possible. However, it may be the case that there also are some propositions that cannot figure in contradictory thoughts such as (3). That is, there may be consistency constraints on thoughts that render impossible certain kinds of thought-contradictions.

Now, suppose that P is a proposition that cannot figure in a thought that "P and ¬P." That is, suppose that (3) is an impossible thought. If (3) is impossible, and if (1) and (2) entail (3), then it follows that it's not the case simultaneously that (1) and (2). But (1) and (2) *can* be true simultaneously. So it must be false that (1) and (2) entail (3). And the best explanation of how the entailment could fail is if (1) and (2) belong to different thought-worlds.

This is an alternative, Williams-inspired argument for the necessity of relativizing (1) and (2) to their own thought worlds. The argument relies on the following Multiplicity Principle for thought-worlds:

> **Thought Consistency Multiplicity Principle (TCMP)**: Necessarily, for any pair of thoughts T_1 and T_2 with content P and $\neg P$, if (i) T_1 and T_2 are simultaneous, and (ii) P is a proposition of type C, then T_1 belongs to thought-world w_1, and T_2 belongs to thought-world w_2, where $w_1 \neq w_2$.

That P is a proposition of type C means that P is a proposition of such a kind that it's impossible for there to be a thought that P and $\neg P$.

This is a Multiplicity Principle applying at the level of the contents of thoughts. For the vehicles of thought are candidates neither for consistency nor for inconsistency.

3.3 A Experiential Consistency Multiplicity Principle

Presenting the Principle

Taking inspiration from TCMP, it's possible to formulate a non-contradiction Multiplicity Principle for experiences:

> **Experiential Consistency Multiplicity Principle (ECMP)**: Necessarily, for any two experiences e_1 and e_2 with contents c_1 and c_2, if (i) e_1 and e_2 are contemporaneous, and (ii) c_1 is inconsistent with c_2, then e_1 belongs to stream of consciousness E_1, and e_2 belongs to stream of consciousness E_2, where $E_1 \neq E_2$.

ECMP is the claim that *experiential contradictions* must belong to different streams of consciousness. An experiential contradiction is a pair of experiences e_1 and e_2 whose contents are inconsistent. ECMP is a principle about the contents of experience because vehicles of experience aren't candidates for consistency or inconsistency.

An example of an experiential contradiction would be an experience e_1 of an apple as being uniformly red and a simultaneous experience e_2 of the same apple as being uniformly green. Assuming that uniform redness and uniform greenness are incompatible qualities, e_1 and e_2 have inconsistent contents and are thus contradictory.

Admittedly, the contents of e_1 and e_2 aren't *logically* contradictory. For it's not logically contradictory for an apple to simultaneously be uniformly green and uniformly red, say. But the contents are contradictory in the non-technical sense of its being impossible and inconceivable for an apple to simultaneously be red all over and green all over.

Experiential contradictoriness shouldn't be conflated with *experiential indeterminacy*. An experience is indeterminate if there's an indeterminacy in its content. For example, an experience *e* of an apple is experientially indeterminate if there's no fact of the matter as to whether *e* presents the apple as uniformly red or uniformly green. In contrast, an experiential contradiction doesn't involve indeterminacies of content, but rather conflicting content-level determinacies. Instead of being constituted by a single experience with indeterminate content, an experiential contradiction is constituted by two experiences with inconsistent content.

It bears remarking that ECMP applies to contemporaneous experiences. In order for there to be a contradiction between the contents of e_1 and e_2, the experiences must present the world as being a particular way and simultaneously as being another, incompatible way. If e_1 and e_2 aren't contemporaneous, their contents will just be different rather than inconsistent. An apple's being red at t_1 and green at t_2 *isn't* contradictory.

Defending the Principle

In this section, an argument for the truth of ECMP is presented. The argument will take as its starting point the claim that consciousness is *transparent* (henceforth "Transparency"). There are various ways that transparency claims are used in the philosophy

of mind. However, for our purposes here, Transparency will be understood as the claim that the contents of a subject's experiences are presented to her as ways the world is. Experience has the character of presenting its objects, qualities, and relations as states of affairs in the external, mind-independent world. From the standpoint of how experience appears to us, it's a window into a world existing outside us.

Note that Transparency, as characterized here, isn't the claim that the contents of a subject's experiences are ways the world is. For it's at least conceivable that there are situations, most notably hallucinations, in which the contents of experience don't reflect how the world is. Instead, Transparency only requires that the contents of a subject's experiences are *presented* to her as ways the world is. And experience-contents can be presented in this way even when they don't correspond to how the world is.

On the basis of Transparency, the following argument for ECMP can be formulated:

P1: For any state of affairs A, if A is impossible, then there cannot be an experience of A.

P2: Contradictions are impossible.

C: Therefore, there cannot be experiences of contradictions.

P1 follows from Transparency. According to Transparency, an experience of A must present A as a way the world is. And A can be presented as a way the world is only if A

is possible. Therefore, if *A* is impossible and Transitivity holds, then there cannot be an experience of *A*.

P2, in turn, is trivial. Contradictions are impossible. Hence, experiences of contradictions are experiences of impossible states of affairs. Given P1, there cannot be experiences of impossible states of affairs. Thus, experiences of contradictions are impossible (C). And if it's impossible to experience contradictions, then simultaneous experiences with inconsistent content must belong to different streams of consciousness.

An objection to this argument has to target P1, as P2 is obviously true. To show that P1 is false, one must show that it's possible to experience impossible states of affairs. There are some fairly good candidates for being experiences of impossible entities. For example, the lithographs and woodcuts of M.C. Escher contain many examples of impossible objects such as infinitely ascending and descending stairs.[12] One can argue that, when one looks at an Escher-style depiction of impossible objects, one has an experience of impossible states of affairs. This would constitute a counterexample to P1.

At least two different things can be said in response. The first one is to deny that one has an experience of an impossible object when looking at Escher's lithographs and woodcuts. Instead, on this view, one has an experience of a plurality of possible objects (such as the different sections of the stairs) which are confused for an impossible object.

The second response is to acknowledge that experiences of Escher's artworks have impossible entities as their objects. By doing so, one concedes that P1 isn't unqualifiedly

[12] Many of these can be viewed on the website of the M.S. Escher Foundation and the M.S. Escher Company: https://mcescher.com/.

true. However, that's still consistent with holding that *some* impossible states of affairs are such that they cannot be experienced. And contradictions in the sense defined above are good candidates. Even though it may be possible to experience one of Escher's infinitely ascending and descending stairs, that doesn't imply that it's possible to simultaneously experience an apple as uniformly red and as uniformly green.

Thus, even though the objection to P1 is a challenging one, there are multiple ways to get around it. I'll henceforth assume that the argument for ECMP is sound and valid.

4 Forms of Unity

This section describes four different kinds of Unity Relation: *Spatial Unity*, *Temporal Unity*, *Phenomenal Unity*, and *Access Unity*. As noted in section 2, Unity SSPs and Unity NSPs construe conscious unity as either sufficient or necessary for a set of experiences to belong to a stream of consciousness. Thus, there's a form of Unity SSP and Unity NSP corresponding to each form of conscious unity described here.

Only *synchronic* forms of conscious unity are discussed in what follows. These are unity relations between simultaneous experiences. There also may be *diachronic* unity relations, obtaining between non-simultaneous experiences. However, diachronic unity relations won't be discussed in this thesis.

4.1 Spatial Unity

Experiences can be unified in virtue of being spatially related at the level of the contents of experience. This principle can be formulated as follows:

Spatial Unity: Necessarily, for any experiences e_1 with content c_1 and e_2 with content c_2, e_1 and e_2 are spatially unified iff there's an experience E of a spatial relation between c_1 and c_2.

Spatial Unity is a form of unity concerned with spatial relations holding at the level of contents of experience. It captures the sense in which the objects and qualities encountered by a subject in experience are presented as being spatially related. For example, I'm currently having an experience e_1 of my laptop and an experience e_2 of my notebook. The *contents* of e_1 and e_2 stand in a spatial relation; currently, my notebook is placed to the right of my laptop. Conversely, there's no reason why the *vehicles* of e_1 and e_2 should stand in a corresponding spatial relation.

In order for the contents of e_1 and e_2, c_1 and c_2, to be spatially related, the spatial relation between the two must itself be a content of experience. Since there cannot be a content of experience without an experience, it follows that there's an experience E of the spatial relation between c_1 and c_2. The object of E is the content-level spatial relation between c_1 and c_2.

Ordinary experience is replete with examples of experiences' being spatially unified, of contents of experience being presented as spatially related. But what does it take for experiences to *not* be spatially unified?

An example of this could be a case in which a subject's visual field is split in two halves, with the contents of experience within each half being presented as spatially related, but with no presentations of spatial relations between contents corresponding

to *different* halves of the visual field. A subject facing this situation would be experientially presented with two spatially well-defined regions of reality, but the area "between" these regions wouldn't be spatially well-defined. An example would perhaps be horses, since they have two visual fields due to the separation of their eyes.

4.2 Temporal Unity

The next form of unity under consideration concerns temporal relations between contents of experience. Specifically, it is a principle about how experiences can be unified by having contents that are presented to a subject as occurring *simultaneously*.

> **Temporal Unity**: Necessarily, for any experiences e_1 with content c_1 and e_2 with content c_2, e_1 and e_2 are temporally unified iff there's an experience E of a simultaneity relation between c_1 and c_2.

Temporal Unity is a synchronic form of unity concerned with simultaneity relations holding at the level of contents of experience. This conception of unity reflects how the events and qualities that a subject experiences are presented to her as standing in simultaneity relations—how they're presented as occurring contemporaneously in her consciousness. For example, I'm currently having an experience e_1 of a red apple and an experience e_2 of squawking Canada Geese. The *contents* of e_1 and e_2 stand in a simultaneity relation; they're experientially presented to me as occurring together, simultaneously.

Given its focus on simultaneity, Temporal Unity is a synchronic relation. But there may well be other, diachronic forms of temporal unity relations. They won't be discussed here, however.

In order for the contents of e_1 and e_2, c_1 and c_2, to stand in a simultaneity relation, the simultaneity relation between the two must itself be a content of experience. Since there cannot be a content of experience without an experience, it follows that there's an experience E of the simultaneity relation between c_1 and c_2. The object of E is the content-level simultaneity relation between c_1 and c_2.

What would be a case in which a subject's experiences *aren't* temporally unified? One answer could be that temporally disunified experiences have vehicles that don't occur in temporal proximity to each other. For example, the vehicle of my latest experience of drinking wine occurred yesterday at 9 PM, while the vehicle of my latest coffee experience occurred today at 8 PM. Obviously, due to the temporal separation of the vehicles of the wine-experience and the coffee-experience, the contents of those experiences cannot be experientially presented to a subject as simultaneous.

However, this doesn't get to the core of what a *failure* of Temporal Unity is supposed to be. For two experiences with temporally distant vehicles aren't even candidates of being temporally unified in the first place. Thus, they can hardly be construed as *counterexamples* to consciousness' being temporally unified. Specifying what a failure of Temporal Unity is involves a specification of experiences that are eligible candidates for standing in the relation. Note that it's not necessary to take a stance here on whether it's possible for eligible candidates not to stand in the relation. Here, the concern is

simply specifying the hypothetical circumstances under which eligible candidate experiences would fail to be temporally unified.

A potentially more promising specification of such circumstances is one according to which e_1 and e_2 are temporally disunified iff (i) the experiences' vehicles, v_1 and v_2, occur simultaneously, and (ii) the experiences' contents, c_1 and c_2, aren't experientially presented as simultaneous. That, at least, would be a specific condition of temporal disunity. However, there's no a priori guarantee that vehicular simultaneity necessarily coincides with content simultaneity. At the very least, it's extremely empirically presumptuous to assume that vehicular non-simultaneity is necessary and sufficient for temporal disunity.

Thus, it's extremely difficult to sharply delineate the conditions under which experiences that are eligible candidates for Temporal Unity are temporally disunified. Instead, we'll have to make do with the following unsatisfyingly loose condition: That two experiences e_1 and e_2 are temporally unified only if the vehicles of e_1 and e_2, v_1 and v_2, occur "close" to each other in a short temporal interval. That rules out experiences such as the wine-experience and the coffee-experience from being eligible candidates of being temporally unified. However, this is just a provisional surrogate for a more perspicuous elucidation of the conditions under which experiences fail to be temporally unified.

4.3 Phenomenal Unity

Phenomenal Unity is a form of synchronic unity relation obtaining between phenomenally conscious experiences. A mental state is phenomenally conscious iff there's

something it's like for a subject to have that experience (Block, 1995). Accordingly, the notion of Phenomenal Unity seeks to capture how there's something that it's like for a subject to have multiple phenomenally conscious experiences *together*. For example, there's something it's like for me to simultaneously have a visual experience of words appearing in this word document *and* an auditory experience of the clicking of my keyboard.

Phenomenal Unity is sometimes understood in mereological terms. On this view, which may have its origin in Lockwood (1989: 88), two experiences e_1 and e_2 are phenomenally unified iff they're parts of the same experience *E*. The most systematic attempt to articulate a mereological conception of mereological unity is found in Tim Bayne's (2010) book *The Unity of Consciousness* (although see also Bayne, 2008; Bayne & Chalmers, 2003). In the book, Bayne gives the following mereological characterization of Phenomenal Unity:

> We might say that two conscious states are phenomenally unified when, and only when, they are co-subsumed. What it is to experience a headache and the sound of a trumpet together—what it is for these two experiences to possess a 'conjoint phenomenal character'—is for there to be a single experience that in some way includes both the experience of the headache and that of the trumpet [---]. How should we think of subsumption? It is tempting to think of it in mereological terms—that is, in terms of parts and wholes. What it is for one experience to subsume another is for the former to contain the latter as a part [---]. One's overall phenomenal field is an experience that contains withing itself experiences, nestled like Russian dolls within each other. (Bayne, 2010: 20-21.)

Bayne writes that two experiences e_1 and e_2 are phenomenally unified iff they're "co-subsumed," and that e_1 and e_2 are co-subsumed iff they're part of the same experience *E*. This principle can be formulated as follows:

Mereological Phenomenal Unity (MPU): Necessarily, for any two simultaneous phenomenally conscious experiences e_1 and e_2, e_1 and e_2 are phenomenally unified iff there's a phenomenally conscious experience E of which e_1 and e_2 are parts.

Interestingly, MPU describes a relation between the *vehicles* of experience. For, according to the principle, the mark of phenomenality is standing in certain mereological relations. And only vehicles of experience—experiences construed as spatiotemporal events in the world—can be intelligibly understood as standing in mereological relations. Bayne uses a metaphor of spatial containment (Russian dolls) to elucidate the kind of mereological relations he has in mind. Contents of experience—ways that the world is presented to an experiencing subject—cannot stand in such containment relations.

Admittedly, even if Phenomenal Unity relations are mereological relations between experience-vehicles, those relations could still reflect relations of "shared what-it's-likeness"—relations of "conjoint phenomenal character," to use Bayne's locution (2010: 20.) These will be relations at the level of the contents of experience, at the level of how the world is presented in consciousness. Aren't these the *real* Phenomenal Unity relations? Phenomenal Unity is better understood in terms of a content-characterization than in terms of a vehicle-characterization.

A way of characterizing Phenomenal Unity as relation between the contents of experience is as follows:

Content-Level Phenomenal Unity (CPU): Necessarily, for any two simultaneous phenomenally conscious experience e_1 and e_2 with contents c_1 and c_2 and subject *S*, e_1 and e_2 are phenomenally unified iff *S* experiences c_1 *and* c_2.

Let's explain CPU in terms of an example. Suppose *S* simultaneously has a visual experience e_1 of a red tomato and a gustatory-olfactory experience e_2 of cinnamon. According to CPU, e_1 and e_2 are phenomenally unified iff it's the case that S experiences the tomato *and* the cinnamon. In contrast, e_1 and e_2 aren't phenomenally unified if it's the case that *S* experiences the tomato while *separately* experiencing the cinnamon.

The difference between Phenomenal Unity and lack of Phenomenal Unity is expressed by a scope distinction, a distinction between "*S* experiences (c_1 and c_2)" and "*S* experiences (c_1) and *S* experiences (c_2)." This scope distinction captures the difference between, on the one hand, there being a single presentation of two contents of experience for a subject, and, on the other hand, there being two presentations of content to the subject. Likewise, the scope distinction captures the difference between what it's like to have two phenomenally unified experiences and two experiences that aren't phenomenally unified. In the former case, there's something it's like for *S* to experience the tomato and the cinnamon together, while in the latter case, there's just something it's like for *S* to experience the tomato and something it's like for *S* to experience the cinnamon.

Since CPU is a content-construal of Phenomenal Unity, it doesn't have any implications for mereological relations between the vehicles of experience, unlike MPU. In

what follows, the form of Phenomenal Unity that will be discussed is CPU rather than MPU.

4.4 Access Unity

Access Unity is a form of unity relation obtaining between access conscious mental states. A mental state is access conscious iff it's poised for rational control of action, speech, and inference by a subject (Block, 1995). Accordingly, two experiences e_1 and e_2 are access unified iff they're *jointly poised* for fulfilling this role:

> **Access Unity**: Necessarily, for any two simultaneous experiences e_1 and e_2 with a subject *S*, e_1 and e_2 are (synchronically) access unified iff e_1 and e_2 are jointly available to *S* for rational control of *S*'s action, inference, and speech.[13]

Access Unity is best construed as a relation between the vehicles of experience. For if e_1 and e_2 are access unified, then they stand in a relation facilitating certain forms of behaviour on the part of the subject. In other words, it's a relation between experiences themselves rather than a relation between their contents. However, the obtaining of Access Unity relations between experiences mirror the role that the contents of such experiences play in reasoning.

Throughout the thesis, Access Unity will mainly play the role of being a behavioural test of the presence or absence of other kinds of unity relations between

[13] For similar definitions of Access Unity, see (Bayne & Chalmers, 2003: 29 ; Schechter, 2018: 26.)

experiences. That is, if two experiences e_1 and e_2 have vehicles that are access conscious, that constitutes good evidence that the contents of e_1 and e_2 stand in content-level unity relations, such as Spatial Unity, Temporal Unity, and Phenomenal Unity. Conversely, if e_1 and e_2 aren't access conscious, that's an indication that the contents of the experiences don't stand in content-level unity relations.

4.5 The Relation between the Spatiotemporal Unity and Phenomenal Unity

What's the relation between Spatial Unity, Temporal Unity, and Phenomenal Unity?

It's clear that experiences that are spatially or temporally unified must be phenomenally unified. Recall from the definition of Spatial and Temporal Unity that two experiences e_1 and e_2 are spatially, or temporally, unified iff there's an experience E of a spatial, or temporal, relation between the contents c_1 and c_2 of e_1 and e_2. If a subject S has an experience of the spatiotemporal relation between c_1 and c_2, then S doesn't experience c_1 while separately experiencing c_2: S experiences c_1 *and* c_2.[14] Thus, there's something it's like for S to have e_1 and e_2 together, meaning that the experiences are phenomenally unified.

Must experiences that are phenomenally unified be spatially or temporally unified? When it comes to Spatial Unity, the answer is *no*. Just because S experiences c_1 and c_2 together doesn't mean that they are experientially presented to S as standing in a spatial relation. For example, my experiences of backpain and of hunger are phenomenally unified, but their contents aren't presented to me as spatially related.

[14] In the context of discussing unity relations, the word "spatiotemporal" is used in the sense of "spatial or temporal."

The answer to the question whether there can be Phenomenal Unity without temporal unity is more complicated. For Phenomenal Unity is normally presented as a *synchronic* unity relation: A relation requiring the simultaneity of its relata. The relevant notion of simultaneity is a content-conception rather than a vehicle-conception. In order for two contents to be experienced as occurring together, as having a shared what-it's-likeness, the contents must be experientially presented as simultaneous. In contrast, two experience-vehicles' occurring simultaneously doesn't guarantee that the vehicles have contents that are experientially presented as occurring *together*.

Provided that the pertinent simultaneity relation between phenomenally unified experiences cannot be vehicular simultaneity, the only way to reject the claim that Phenomenal Unity entails Temporal Unity is by removing the requirement that phenomenally unified experiences are simultaneous. However, it's difficult to conceive what Phenomenal Unity in the absence of Temporal Unity would amount to. Perhaps such a notion is unintelligible.

If the foregoing is correct, then Temporal Unity would be both sufficient *and* necessary for Phenomenal Unity. And the best explanation of this may be the *identity* of the two relations. Nevertheless, in what follows, Phenomenal Unity and Temporal Unity will be discussed as distinct notion.

A brief final point should be made about the relationship between Spatial Unity and Temporal Unity. The notion of Spatial Unity that's of concern in this thesis is *synchronic* Spatial Unity: Two experience-contents' being experientially presented as simultaneous *at a time*. Since the relevant notion of "time" is a content-notion, spatially

unified experiences must be temporally unified; their contents must be experientially presented to a subject as occurring simultaneously.

5 Splitting the Brain

SB-patients have undergone commissurotomy or callosotomy surgery. In those surgical procedures, the tract of nerve fibres connecting the two cerebral hemispheres (the corpus callosum) is surgically severed. Commissurotomy and callosotomy have been performed since 1930s on patients suffering from severe epilepsy. Specifically, the purpose of the procedure was to stop the interhemispheric propagation of epileptic fits. The first surgeries of this kind didn't considerably alleviate the severity of the patients' epileptic fits (van Wagenen & Herren, 1940). Subsequent check-ups on the patients did reveal, however, that the surgery hadn't significantly impacted on the patients' behaviour in any systematic way (Akelaitis, 1941, 1944). This prompted speculations that the corpus callosum didn't have a significant role to play in physiological architecture of the brain. Famously, it was even joked that the sole purpose of the corpus callosum was the interhemispheric transmission of epileptic fits (McCulloch, 1949).

Following the lukewarm success of early experiences with split-brain surgery, there was a long lull in research on SB-patients. However, this changed during the 1960s and 1970s due to research by Roger Sperry, Michael Gazzaniga, and Joseph Bogen (among others). This period saw the emergence of more effective methods for commissurotomy and callosotomy surgery for epilepsy (Bogen & Vogel, 1962). Moreover, psychological experiments conducted on SB-patients, what I'll call "SB-experiments," sparked an explosion of scientific interest in brain organization and the significance of the corpus callosum for human cognition. Notable experiments were conducted by Sperry and

Gazzaniga (see, e.g., Gazzaniga et al., 1965; Gazzaniga & Hillyard, 1971; Gazzaniga & Sperry, 1967; Sperry, 1968, 1974; Sperry et al., 1969). These experiment allowed neuropsychologists to study how the cerebral hemispheres functioned without callosal means of information transfer. The experiments showed that, under carefully controlled experimental conditions, there can be breakdowns of interhemispheric informational integration in SB-patients. In these experiments, the hemispheres were unable to transmit information to each other. This was manifested in the patients' exhibiting *prima facie* odd forms of behaviour, such as denying having seeing an object in their left visual fields even though the patients could draw or manually identify the object using their left hands (Sperry, 1968).

In SB-experiments, the two hemispheres become partially autonomous systems for processing information and producing behaviour. As Gazzaniga writes in a retrospective summary of the research landscape since he started conducting research on SB-patients in the 1960s…

> [severing] the entire callosum blocks the interhemispheric transfer of perceptual, sensory, motor, gnostic and other forms of information in a dramatic way, allowing us to gain insights into hemispheric differences as well as the mechanisms through which the two hemispheres interact […]. (Gazzaniga, 2005: 653-654.)

SB-experiments have been important for advancing knowledge of the cognitive structure of the brain. In particular, SB-experiments have provided neuropsychologists with new possibilities for studying the specialization of the hemispheres for different kinds of cognitive-behavioural tasks. For example, much of what is known about the left LH's specialization for language processing and RH's specialization for spatial reasoning can

be traced back to experiments on SB-patients (see, e.g., Gazzaniga & Sperry, 1967; Sperry, 1974).

In everyday life, SB-patients are often indistinguishable from neurotypical individuals (i.e., individuals with a corpus callosum), and are able to perform most basic cognitive-behavioural functions satisfactorily. Sperry reports the following:

> A person two years recovered from the operation and otherwise without complications might easily go through a routine medical check-up without revealing that anything was particularly wrong to someone not acquainted with his surgical history. Speech, verbal intelligence, calculation, established motor coordination, verbal reasoning and recall, personality and temperament are all preserved to a surprising degree in the absence of hemispheric interconnection. (Sperry, 1974: 6.)

Nevertheless, despite the normality of SB-patients under extra-experimental conditions, their behaviour in SB-experiments soon prompted speculations that the structure of their consciousness was unusual. For example, in one of Sperry's papers, he states that SB-experiments show that SB-patients have multiple streams of consciousness (cf. Sperry, 1974: 7), one corresponding to each hemisphere:

> [---] [There's] an apparent doubling in most of the realms of conscious awareness. Instead of the normally unified single stream of consciousness, these patients behave in many ways as if they have two independent streams of conscious awareness, one in each hemisphere, each of which is cut off from and out of contact with the mental experience of the other. In other words, each hemisphere seems to have its own separate and private sensations; its own perceptions [...]. (Sperry, 1968: 724.)

Later in the paper, Sperry goes even further, seemingly suggesting that the hemispheres not only be thought of as loci of separate streams of consciousness, but also as constituting distinct, fully separate conscious minds:

> [...] [I] favor the view that in the minor hemisphere we deal with a second conscious entity that is characteristically human and runs along in parallel with the more dominant stream of consciousness in the major hemisphere. (Sperry, 1968: 732.)

Philosophers, in turn, have been quick to follow neuropsychologists in interpreting SB-consciousness as having a structure that's at least *prima facie* puzzling. In a 1971 paper about the philosophical implications of new research on SB-patients, Thomas Nagel writes that SB-experiments profoundly challenge conception of SB-patients as being single conscious entities:

> The feature of the mentalist conception of persons which may be recalcitrant to integration with these data is not a trivial or peripheral one, that might easily be abandoned. It is the idea of a *single* person, a single subject of experience and action, that is in difficulties. (Nagel, 1971: 396.)

What has been thought to motivate these conclusions about the structure of SB-patients' consciousness are certain controversial SB-experiments, mainly conducted by Sperry and Gazzaniga in the 1960s and 1970s. Let's discuss these experiments.

6 Experiments on Split-Brain Patients

6.1 How Split-Brain Experiments Work

SB-experiments are intended to test how surgically disconnected hemispheres process information and execute behavioural tasks when forced to rely strictly on direct (intracranial) means of interhemispheric integration and transmission of information. In other words, the purpose of SB-experiments is to elucidate the extent to which cerebral hemispheres can integrate and transmit information without the corpus callosum and without reliance behavioural intermediaries, such as movement of patients' eyes and hands. Thus, the challenge of designing a SB-experiment lies in concocting ways of isolating the information received by each hemisphere, so that it can be assessed how much of that information is transmitted to the other hemisphere.

An important way in which neuropsychologist attempt to meet this challenge is by exploiting the *lateralization* of sensory stimuli.[15] Sensory stimuli are lateralized in the sense that different kinds of stimuli are processed in different "sides" of the perceptual system, and information from the different sides are often transmitted to different hemispheres. For example, visual stimuli are transmitted *contralaterally*. Visual information on stimuli in the right half of the visual field (RVF) is transmitted to LH. Correspondingly, visual information from the left half of the visual field (LVF) is transmitted to RH. Tactile information is also transmitted and processed contralaterally. Tactile information about things felt by the right hand is transmitted to LH, while tactile information about things felt by the left hand is transmitted to RH.

Related to lateralization, SB-experiments are also often built around the hemispheres' specialization for executing different behavioural and cognitive functions. One example is that, by and large, hand movements are controlled contralaterally. Another, and perhaps even more important, point is that, by and large, RH cannot produce speech. As a consequence, SB-patients can verbally report RH-associated sensory information only if that information has been transmitted to LH. One of the main lessons of SB-experiments is that there are certain, highly specific contexts in which interhemispheric transmission of information is impossible without the corpus callosum.

In the traditional method for conducting SB-experiments, sensory stimuli were lateralized using a machine called a *tachistoscope*. For an overview of how such

[15] That isn't to say that all forms of sensory stimuli are unilateralized. For example, information on auditory stimuli is transmitted to both LH and RH, although not to the same extent.

experiments work, see, e.g., (Gazzaniga et al., 1965; Sperry, 1974; Sperry et al., 1969). The focus will mainly be on visual perception here. The SB-patient is seated in front of two white screens, one to either side of a fixed point that the patient is instructed to gaze at fixedly. These screens also serve the purpose of obscuring the patient's hands. Visual stimuli about objects to the right of the focal point is transmitted to LH, while stimuli about objects to the left is transmitted to RH. Images are then flashed for one tenth of a second on either screen. The short screening-time means that the patient doesn't have time to move her eyes and thus expand the range of visual information that's available to each hemisphere. Experimenters have thereby created a testing situation in which RH-associated information is available to LH only if interhemispheric information transmission is possible.

For neurotypical subjects, interhemispheric information transmission is achieved through the corpus callosum. However, that mechanism isn't available to SB-patients. Consequently, if SB-patients *are* capable of systematically transmitting information interhemispherically in experimental conditions, there are two potential explanations.

Firstly, there could be *subcortical* means of integrating information across the hemispheres. This would mean that brain structures located beneath the corpus callosum, structures are more evolutionarily ancient than the cortical commissures, could ensure a degree of interhemispheric integration of information. This suggestion matches what one finds in certain SB-experiments, where affective states such as arousal have been observed to be interhemispherically transmitted.[16] However, the majority view in

[16] For, example, in one experiment, a SB-patient was tachistoscopically presented with a photograph of a naked woman in the patient's LVF. The visual information was transmitted to RH. Nevertheless, using LH, the patient was able to remark that they felt aroused and embarrassed, even though they couldn't

neuropsychology is that most forms of information, including sensory information, cannot be transmitted interhemispherically through subcortical mechanisms (Gazzaniga & Hillyard, 1971; Kingstone & Gazzaniga, 1995; Seymour et al., 1994). And the concern of this chapter is specifically with the interhemispheric integration of *sensory* information.

Secondly, instances of integration of sensory information in SB-experiments may be due to so-called *behavioural cross-cuing* between the hemispheres. These are extracranial and behavioural strategies that SB-patients often utilize to integrate information across the hemispheres in the absence of a direct callosal pathway between them. We've already mentioned one such integration strategy, namely moving one's eyes. However, there are other unwitting integration strategies that are remarkably ingenious. For example, in one experiment (Sperry et al., 1979), a SB-patient was asked to name the word "cat," which had been tachistoscopically presented in LVF, and hence transmitted to RH. When asked to name the object (using LH), the patient used her left hand (associated with RH) to sketch "cat" on her right hand. Since the right hand was in RVF, the patient was able to name the object with her LH. This example shows that eliminating the possibility of behavioural cross-cuing in SB-experiments is a far from trivial matter; indeed, for certain experiments, it has been deemed impossible (Seymour et al., 1994).

specify the source of those feelings. Consequently, affective information was transferred from the RH to LH even though the visual information wasn't interhemispherically shared (Sperry, 1968: 732).

6.2 Two Split-Brain Experiments

In this section, two examples of SB-experiments are discussed. Although there's a plethora of kinds of SB-experiments, discussing these two will illustrate the core issues about SB-patients that are interesting from a philosophical point of view.

First Experiment

The first SB-experiment that will be discussed here is commonly recurring in papers on SB-patients. A SB-patient is asked to perform an inter-modal object recognition task involving visual perception and stereognosis (tactually perceiving the form of objects). In the experiment (Gazzaniga & Sperry, 1967: 139-140), the patient was tachistoscopically presented with the word "comb" in LVF but not in RVF. When the patient was asked what object they saw, they responded that they didn't see any object (since LH only receives information from RVF). However, when asked to use their left hand in an attempt to retrieve the object they denied to have seen, the patient was able to stereognostically identify the comb from a range of other alternatives. When asked what object they had selected, the patient issued a random guess.

What seems clear about this experiment is that sensory stimuli aren't interhemispherically integrated even though RH can recognize and respond to the stimuli in complex ways. Interhemispherically, LH doesn't receive visual information from RH about the word flashed in the SB-patient's LVF. And it's likely that RH receives visual information from LVF. For, unless RH registered the visual stimulus, there wouldn't be an explanation of the patient's capacity to stereognostically identify the comb using her RH-controlled left hand. It's also noteworthy how advanced this object recognition

task is. It involves, firstly, RH receiving visual stimuli from the LVF; secondly, RH processing and comprehending the semantic information encoded by the visual stimuli (i.e., the word "comb"); thirdly, RH retaining in short term memory information about the semantically relevant objects; fourthly, RH recognizing what tactile properties an object of the semantically relevant kind would have (i.e., a comb); and, fifthly, RH stereognostically discriminating the object from an array of other candidates. As Sperry writes (about a similar experiment with the same structure), RH performs actions with a degree of cognitive complexity that hasn't been encountered in non-human animals:

> This perceptual or mnemonic transfer from one sense modality to another has special theoretical interest in that it is something that is extremely difficult or impossible for the monkey brain. The right hemisphere, in other words, may be animallike in not being able to talk or write, but in performances like the foregoing and in a number of other respects it shows mental capacities that are definitely human. (Sperry, 1968: 731.)

What the experiment shows is that RH can be responsible for complex visuo-tactile information integration and cognitively complex behaviour even though RH-associated sensory stimuli isn't transmitted interhemispherically.

Second Experiment

The second SB-experiment discussed in this section is also about vision. Here, though, the focus is on object representation rather than on object recognition. In the experiment (Sperry, 1968: 725-726), two visual stimuli are tachistoscopically presented in one visual hemifield each. In LVF, the patient is presented with a dollar sign ("$"), while in RVF, they're presented with a question mark ("?"). When asked, the patient verbally asserts that they see a question mark and deny that they see any other object. However, when the patient is instructed to draw the object they see using their left

hand and without visual aid, they draw a dollar sign. If subsequently asked about what they have drawn, the SB-patient will respond that they have drawn a question mark.

Just like the previous experiment, this experiment betrays a lack of interhemispheric integration of visual information in SB-patients, even though RH can recognize and respond to visual stimuli in complex ways. LH doesn't receive visual information from RH about the dollar sign flashed in the patient's LVF. And it's likely that RH receives visual information from LVF. For, unless RH registered the visual stimulus, there wouldn't be an explanation of the patient's capacity to draw the dollar sign using her left hand. It's also noteworthy how advanced this object recognition task is. It involves, firstly, RH receiving visual stimuli from LVF; secondly, RH retaining visual information in short term memory for use in a manual task requiring great precision; and, thirdly, RH executing the manual task.

What's noteworthy about the experiment, even compared to the last one, is the *extent* of the disconnect between the hemispheres. Not only is visual information not transmitted from RH to LH at the start of the experiment, but the linguistic utterances produced by LH-associated speech production mechanisms fail to match the behaviour of the patient's left hand. Sperry puts the point quite aptly when he writes that

> […] the one hemisphere does not know what the other hemisphere has been doing. The left and the right halves of the visual field seem to be perceived quite separately in each hemisphere with little or no cross-influence. (Sperry, 1968: 726.)

Summary of Considerations of Experiments

SB-experiments show the possibility of breakdowns in interhemispheric integration of sensory information. Additionally, they show that, at least in specific experimental

situations, the hemispheres function as separate systems for processing sensory stimuli and for producing behaviour.

6.3 Split-Brain Patients and the Hemispheres' Capacity for Experience

I'll assume in what follows that both LH and RH support, or are able to support, consciousness in SB-patients. That is, both cerebral hemispheres can serve as the locus of consciousness.

LH is undeniably capable of supporting experiences. In SB-experiments, patients can utilize LH-controlled speech mechanisms to report the presence of lateralized sensory stimuli (in the case of visual stimuli, in RVF). A slightly more controversial question is whether RH supports consciousness in SB-patients. RH is mainly incapable of producing speech, and in the early reception of research on SB-patients, this was sometimes taken as a sign of the absence of consciousness in RH [insert citation].

However, there are at least two reasons why RH has as good a claim to consciousness as LH. To begin with, patients who have undergone left hemispherectomy, which is the surgical removal of LH, typically behave in a characteristically consciousness-suggesting way. In fact, they even exhibit limited capacities of language production, albeit at a very rudimentary level [insert citations for all of this]. This would be difficult to explain if patients with surgically removed LHs weren't conscious. Thus, there are reasons to believe that RH is able to support consciousness in SB-patients.

Further evidence of this is the complexity of the cognitive behavioural tasks that SB-patients are able to execute. The "comb experiment" presented in the last section shows that SB-patients can use RH to carry out sophisticated object recognition tasks

which require, among other things, the capacity to compare disparate kinds of sensory stimuli (visual stimuli of the word "comb" and tactile stimuli of a comb). If there were no experience of the word or no tactile experience of the comb corresponding to RH, B-patients' ability to execute the object recognition task couldn't be explained. All in all, the behaviour associated with RH is too complex to not be associated with conscious experience. Shallice puts the point nicely when he writes that…

> [if] this level of performance could be obtained unconsciously, then it would be really difficult to argue that consciousness is not an epiphenomenon. Given that it is not, it is therefore very likely, if not unequivocally established, that the split-brain right hemisphere is aware. (Shallice, 1997: 264.)

6.4 Split-Brain Patients and the Malleable Stream View

Now that the foundations of SB-experiments have been described, we're in a better position to describe the theory defended in this theory, what I've called the Malleable Stream View. This is the view that SB-patients have a single stream of consciousness, whose simultaneous constituents are unified extra-experimentally but disunified intra-experimentally. It's possible now to appreciate why this is the case. Outside of SB-experiments, patients can utilize a plethora of extra-cranial mechanisms for interhemispheric integration of experiments. On the Malleable Stream Views, these mechanisms for interhemispheric information integration are precisely the mechanisms that ensure interhemispheric unity between SB-patients' experiences. The breakdown in unity during SB-experiments is precisely due to the unavailability of the unity-

conferring behavioural cross cuing mechanisms that SB-patients have to rely on, given the absence of a direct, intra-cranial mechanism of interhemispheric information transfer.[17]

An interesting implication of this view, if it's correct, is that unity relations between experiences don't have to be grounded exclusively in the neurological architecture of the brain. SB-patients' ability to interact with their environment, most straightforwardly by unconsciously making frequent eye movements, may play an indispensable role in ensuring the unity of their experiences. Consequently, if this is right, then the structure of SB-patients' consciousness isn't something that's intrinsic to them. It's a relational feature, which depends on the environment in which these patients are placed.

[17] In the final stage of editing this thesis, it came to my attention that this view about the unifying nature of behavioural cross-cuing in extra-experimental settings already has been defended in print by Marks (1980). I didn't know this when I formulated the Malleable Stream View.

7 Split-Brain Consciousness and Multiplicity Principles

Preliminaries

One way in which an argument for the Malleable Stream View could stop dead in its track would be if it could be shown that SB-patients' experiences satisfy the antecedent of a Multiplicity Principle. For then it would inevitably follow that SB-patients have multiple streams of consciousness. Which of the three Multiplicity Principles is most relevant in the context of a discussion of SB-patients?

The two main candidates here are Peacocke's Multiplicity Principle and ECMP.[18] To recapitulate, according to Peacocke's Multiplicity Principle, it's impossible for there to be multiple simultaneous experience token of the same maximally determinate kind in a single stream of consciousness. However, even if this is true, it's not really a helpful Multiplicity Principle in the context of discussing SB-patients. For there aren't any well-delineated conditions under which an application of the principle would be warranted. How, precisely, would one go about determining whether two simultaneous experiences of a SB-patients are of *exactly the same* qualitative type? Answers aren't forthcoming.

That leaves ECMP, the principle that there cannot be *experiential contradictions* in a stream of consciousness. That is, if e_1 and e_2 have inconsistent contents, e_1 and e_2

[18] Williams' Multiplicity applies at the level of thought, or belief, or judgement, so it doesn't quite match the question under consideration here, which is how many streams of consciousness SB-patients have (as opposed to the question of how many "thought worlds" they have).

belong to different streams of consciousness. An example of an experiential contradiction would be an experience of one's entire visual field being red, and a simultaneous experience of one's entire visual field being green.

This section of the thesis discusses whether ECMP can be successfully employed to demonstrate that SB-patients have two streams of consciousness. It will be argued that the principle cannot be successfully employed in this way, and that the Malleable Stream View has nothing to fear from ECMP.

A Hypothetical Test of the Consistency of SB-Patients' Experiences

Normally, SB-experiments aren't designed to test philosophical principles about consciousness. In the first instance, they are tests of interhemispheric information integration, hemispheric specialization, and so on. To my knowledge, there aren't any SB-experiments that either directly or indirectly touch on whether it's possible for SB-patients to have contradictory experiences. Therefore, in this subsection, we'll consider a *hypothetical* experiment designed to ascertain whether SB-patients can have contradictory experiences. It will emerge that there are principled philosophical reasons why experiments of the kind presented below cannot establish the possibility of SB-patients' having contradictory experiences.

The hypothetical SB-experiment, which will be called "the Colour Experiment," has four main steps.

In the first step of the experiment, visual stimuli are tachiscopically presented in both RVF and LVF for a subject *S*. Specifically, RVF is uniformly presented to be red,

while LVF is uniformly presented to be green.[19] The colours are flashed with sufficient rapidity to preclude eye movement.

In the second step of the experiment, *S* is asked to specify what colour she experiences. Since speech production is controlled by LH, and LH has access to visual stimuli from RVF but not from LVF, given the sectioning of *S*'s corpus callosum, *S* is likely to reply "red." Furthermore, *S* is likely to respond "no" to the question of whether she experiences any other colour than red.

In the third step of the experiment, *S* is presented with visual stimuli in LVF but not in RVF. Specifically, *S* is visually presented with a photograph of a number of buttons, each with a different colour and each with a distinctive shape (e.g., a circle, a rectangle, an octagon). Included among the buttons are a red button and a green button, as well as some buttons with other colours.

In the fourth step of the experiment, *S* is blindfolded and told the following: "A few moments ago, you were shown a colour, and then a collection of different-coloured and different-shaped buttons. These buttons are now in front of you. Can you use your left hand to press the button whose colour matches the colour you were shown at the beginning of the experiment?" *If S* is able to perform this task (see below for doubts), then *S* will push the green button. For *S*'s left hand is mainly controlled by RH, and

[19] In Sperry and Gazzaniga's early research on SB-patients, the patients were seated some distance in front of the tachistoscope machine (for an illustration, see Sperry, 1974: 8). Hence, such machines couldn't be used to present the entirety of a subject's RVF as being red and the entirety of her LVF as being green. However, more modern tachistoscopes are sometimes placed directly in front of a subject's eyes, similar to the machines utilized during optical eye examinations.

Admittedly, there will always be a limited degree of overlap between the two visual field. That is, there will be a small visual region that's visually accessible from both hemifield. Thus, for the sake of the experiment, let's stipulate that no colour is flashed in this shared visual area. In the ensuing discussion, I'll not mention this complication.

RH has access to visual stimuli from LVF but not from RVF, given the sectioning of *S*'s corpus callosum.

Before discussing the significance of the Colour Experiment and its anticipated results *vis-à-vis* ECMC, it'll be useful to make some remarks on its design. To begin with, it needs to be clarified why the third step of the experiment is needed. Couldn't the experiment be set up so that *S* identifies the green button in LVF using visual perception rather than stereognosis?

The answer is *no*. For if *S* were allowed to visually identify the green button, then there would be time for eye movement. Thus, LH would be able to access the visual stimuli in LVF. Hand-related motor control isn't entirely lateralized; LH can exert some degree of control over a subject's left hand. Hence, if LH has access to the visual stimuli in LVF, then it's possible that LH could make the left hand press red the button and possibly even stop that hand from pressing the green button. For this reason, a photograph of the button must be tachistoscopically presented to LH in the third step of the experiment.

Realistically, the Colour Experiment may be too complex to test what RH-associated experience *S* has in LVF. For the experiment involves a number of cognitively advanced steps. Firstly, RH must be able to retain information about the colour presented in LVF in the first stage of the experiment. Secondly, in the third step of the experiment, RH must be able to quickly register information about the colour and shape of the buttons that are (very rapidly) tachistoscopically presented in LVF. Thirdly, RH must retain this visual information about the buttons' colour and shape. Fourthly, in carrying out the stereognostic identification task, RH must integrate this retained

visual information about the buttons' colour and shape with tactile information about their shape acquired through touching the buttons with *S*'s left hand. It's certainly possible that there are too many cognitively complex intermediary steps in the experiment to ensure that *S* ultimately will press the green button. Perhaps *S* wouldn't perform above chance, or perhaps S would refuse to push a button at all.

However, in the next subsection, it will be shown that even if the Colour Experiment works the way it's intended, the experiment wouldn't show that S has contradictory experiences.

Discussing the Experiment: Experiential Contradictions and Spatial Completeness

Prima facie, one might think that the Colour Experiment contains an example of an experiential contradiction. *S* verbally asserts that she's seeing red, and that this is the only colour she experiences. Hence, one could think that *S* has an experience e_1 of her visual field as being uniformly red. Correspondingly, one might interpret *S*'s left-handed pressing of the green button, and her unwillingness to press the other buttons, as a sign that *S* has an experience e_2 of her visual field as being uniformly green, an experience that's simultaneous with e_1. An experience of one's visual field being uniformly red and an experience of one's visual field as being uniformly green have inconsistent contents; hence, e_1 and e_2 are contradictory. Given an application of ECMP, it follows that e_1 and e_2 don't belong to the same stream of consciousness.

Whether this is the correct way to understand the experiment depends on how one interprets the claim that e_1 and e_2 are experiences of the *entirety* of *S*'s visual field.[20] On one disambiguation, the claim is plainly fallacious. For if by "visual field," one means the range of features of *S*'s environment that *S* can access through visual perception, then e_1 and e_2 aren't experiences of the entirety of *S*'s visual field. For e_1 is elicited on the basis of the reddish visual stimuli in RVF, while e_2 is elicited on the basis of greenish visual stimuli in LVF. That is, neither experience reflects the full range of visual information that's available to *S*. Instead, they only reflect approximately half of it. Since e_1 and e_2 aren't experiences of the entirety of *S*'s visual field, construed in this sense, but rather experiences of RVF and LVF, it's unclear why e_1 and e_2 should be considered experientially contradictory.

However, there's another possible disambiguation of the notion of an "entire visual field." Given this notion, interpreting e_1 and e_2 as contradictory at least isn't obviously fallacious. Let's introduce and explain the disambiguation, and then discuss its application to the Colour Experiment.

On the alternative understanding of "entire visual field," e_1 and e_2 being experiences of S's entire visual field should be understood as the claim that e_1 and e_2 are *spatially complete experiences*. An experience *e* is spatially complete at *t* iff it's the case that the content *c* of *e* is experientially presented to the subject *S* as constituting the entirety of *S*'s visually accessible spatial world at *t*. More simply put, an experience is spatially

[20] As mentioned in a previous footnote, we're ignoring that there's a limited degree of overlap between RVF and LVF.

complete iff its content is presented in consciousness as exhausting the region of the spatial world that is visually presented to the subject.

This is a rather abstract characterization of the notion of spatial completeness. Let's consider some examples. Suppose, for example, that I'm watching a sunset at a beach, the brilliance of the dying sun illuminating the undulating water of an ocean. Let's focus one a specific experience that I have throughout this process, one where the sun is still noticeable above the horizon. In this case, my overall experience *E* of the beach, water, and sun is *spatially complete*. In my consciousness, *E* is presented as the entirety of the spatial world to which I have visual access. If I turn my head, I'll have another spatially complete experience, corresponding to my new situation in the world. In turn, my experience of the sun, my experience of the water, and my experience of the beach will rank as *spatially incomplete*; they aren't presented in experience as exhausting the character of my subjectively experienced spatial world. Instead, they're best understood as constituents of a spatially complete experience, namely *E*.

The notion of a spatially complete experience is quite similar to the notion of an experience that's formed on the basis of visual stimuli from the entirety of a subject's visual field. However, the two notions are distinct. The first notion pertains to the contents of experience; it has to do with whether the contents of one's experiences are *presented* as exhausting the character of one's visually presented spatial world. In contrast, the second notion pertains to whether one's experiences *do* exhaust the range of visual information that one's perceptual system can access.

Of course, in normal cases, a spatially complete experience will be identical with an experience that's formed on the basis of visual stimuli from a subject's entire visual

field. For example, the sunset experience mentioned above is both spatially complete and formed on the basis of visual information drawn from the entire visual field. However, when it comes to SB-patients and the Colour Experiment, we can ask whether the two notions come apart. That is, can e_1 and e_2 be spatially complete, in the sense that their contents are experientially presented to S as *individually* exhausting the character of her visually accessible spatial world, even though each experience is formed on the basis of visual stimuli from just half the visual field?

In order for e_1 and e_2 to be experientially contradictory, the answer to this question must be *yes*. If, at the level of the contents of experience, e_1 presents the entirety of *S*'s visually accessible spatial world as being uniformly red, while e_2 presents the entirety of *S*'s visually accessible spatial world as being uniformly green, then there would be a content-level inconsistency between e_1 and e_2. For it's contradictory, in the sense relevant for ECMP, for *S*'s visually accessible spatial world to be simultaneously presented as uniformly red and as uniformly green. In short, if e_1 and e_2 are understood as spatially complete experiences, it wouldn't be obviously fallacious to interpret the Colour Experiment as showing that it's possible for SB-patients to have contradictory experiences. ECMP would then entail that SB-patients' have at least two streams of consciousness.

However, it's also possible to hold that e_1 and e_2 are spatially *incomplete*. That is, one could hold that the contents of e_1 and e_2 *aren't* presented as individually constituting the totality of *S*'s experientially presented spatial world. Instead, the contents of e_1 and e_2 could be *aspects* of *S*'s experientially presented spatial world. If that's the case, e_1 and e_2 wouldn't have contradictory contents.

For neurotypical subjects, spatially incomplete experiences have contents that are experientially presented as unified. For example, in the sunset experience described above, I experience the sun as being situated *above* the ocean. It's in virtue of standing in such relations that the contents of spatially incomplete experiences can constitute the content of a complete experience, like in the sunset example. However, SB-patients won't have spatially unified experiences, at least in intra-experimental context. As will be shown in section 8 when discussing Spatial Unity and SB-patients, SB-patients have experiences that aren't spatially unified in experiments involving tachistoscopic presentation of visual stimuli.

The problem with construing the Colour Experiment as demonstrating the possibility of SB-patients' having contradictory experiences is that there aren't any good reasons to think that e_1 and e_2 are spatially complete experiences, rather than spatially incomplete experiences that are spatially disunified. The two possibilities are on-par in terms of their ability to account for the (hypothetical) experimental data. Regardless of whether e_1 and e_2 are spatially complete or spatially incomplete, *S* won't linguistically report awareness of anything but the redness presented in RVF. Similarly, if *S* is able to perform all the steps in the experiment, her left hand will only report awareness of the greenness in her LVF (using a button press). It's the breakdown in interhemispheric integration of information which accounts for the behavioural abnormalities that *S* would exhibit in the Colour Experiment.

Hence, the reason for construing e_1 and e_2 as spatially complete experiences and not as spatially incomplete experiences cannot be for empirical reasons. And no *philosophical* reasons for making that assessment are apparent. There just aren't any

satisfactory grounds for holding that e_1 and e_2 must be spatially complete. And without the spatial completeness assumption, the *S* simultaneously having e_1 and e_2 doesn't involve an experiential contradiction. Therefore, ECMP doesn't apply here, and cannot be used to demonstrate that SB-patients have multiple streams of consciousness.

8 Split-Brain Consciousness and Sufficiency Singularity Principles

8.1 The Case against Unity

Unity SSPs and Split-Brain Patients

As a reminder, Unity SSPs are principles according to which a set of experiences $e_1, \ldots, e_n$ belong to a single stream of consciousness if all of $e_1, \ldots, e_n$ are unified with each other. Depending on what version of unity that's at issue, one might end up with Spatial SSP, Temporal SSP, Phenomenal SSP, Access SSP, and perhaps more besides.

One way of arguing for a Unified One-Stream view of SB-patients is by (somehow) establishing that that all contemporaneous experiences of SB-patients are unified. That, in conjunction with an appropriate version of Unity SSP, entails that SB-patients have a single stream of consciousness.

However, this is an untenable strategy for establishing the one-stream view. For it will be shown below that SB-patients have experiences which aren't unified, in any of the senses described in section 4. Thus, although Unity SSPs are true principles, some of SB-patients' experiences don't satisfy their antecedents.

I will begin by making this argument with regard to Spatiotemporal Unity, and then separately make the argument with regard to Phenomenal Unity. Admittedly, discussing the two notions separately may seem superfluous. If Spatiotemporal Unity

entails Phenomenal Unity, then showing that two experiences e_1 and e_2 are phenomenally disunified may be taken as a proof that e_1 and e_2 are spatiotemporally disunified. However, this assumes that the contents of e_1 and e_2, c_1 and c_2, are experientially presented as having a spatiotemporal form. And this doesn't have to be true, at least with regard to space. For c_1 and c_2 can be experientially presented as occurring *together* for a subject without thereby being presented as standing in a spatial relation—consider the backpain-hunger example mentioned in section 4.5. In such a case, e_1 and e_2 being phenomenally disunified wouldn't count as an instance of their being spatially disunified, for e_1 and e_2 wouldn't count as eligible candidates for Spatial Unity in the first place.

Things are a bit trickier when it comes to Temporal Unity. For there are fairly compelling reasons to believe that Temporal Unity is both sufficient *and* necessary for Phenomenal Unity. If that's right, then the contents c_1 and c_2 of any phenomenally unified experiences e_1 and e_2 must be experientially presented to the subject as having a temporal form. Consequently, demonstrating that SB-patients' experiences are temporally disunified would demonstrate their phenomenal disunity. Likewise, given the necessity of Phenomenal Unity for Spatial Unity, e_1 and e_2 being temporally disunified would entail their being spatially unified.

However, even if these entailment relations obtain, it will be instructive to consider the case of Spatiotemporal Unity and Phenomenal Unity separately. For, as will be seen below, an argument for SB-patients' having spatiotemporally disunified experiences must have a slightly different form than an argument for their having phenomenally disunified experiences.

Against Spatiotemporal Unity

Let's focus on the second SB-experiment discussed in section 6.2. In that experiment, a SB-patient *S* is tachistoscopically presented with a dollar sign in LVF and with a question mark in RVF. Suppose that *S* has an experience e_1 of the question mark and an experience e_2 of the dollar sign. What would it mean for e_1 and e_2 to be spatially or temporally unified?

The two experiences would be *spatially unified* iff the question mark and the dollar sign were experientially presented to *S* as standing in a spatial relation. A description of the relation would be that "the dollar sign is to the left of the question mark" or "the question mark is to the right of the dollar sign." Correspondingly, e_1 and e_2 would be *temporally unified* iff the image of the question mark and the image of the dollar sign were presented to *S* as occurring simultaneously. A description of this would be "the question mark and the dollar sign popped up at the same moment," or "I saw the question mark and the dollar sign at the same time."

I'll argue that e_1 and e_2 fulfil neither condition, and hence are spatiotemporally disunified. The broad gist of the argument is that, if e_1 and e_2 were spatiotemporally unified, then *S* would be in a position to specify the spatiotemporal relation between the contents of the experiences. But *S* cannot do this. Hence, e_1 and e_2 are spatiotemporally disunified. The argument will have to take on a slightly different form depending on whether it's used as an argument against the Spatial Unity of e_1 and e_2 or the Temporal Unity of e_1 and e_2. The argument will nevertheless be presented in a general

form. However, the points at which it matters whether it's Spatial Unity or Temporal Unity that's at issue will be identified. Here, in any case, is the argument:

P1: When a SB-patient *S* is simultaneously tachistoscopically presented with visual stimuli in RVF and LVF, then *S* has two experiences e_1 and e_2, with e_1 corresponding to the visual stimuli in RVF and e_2 corresponding with the visual stimuli in LVF.

P2: If e_1 and e_2 are spatiotemporally unified, then *S* is in a position to behaviourally indicate, either verbally or non-verbally, what spatiotemporal relation the contents c_1 of e_1 and c_2 of e_2 stand in.

P3: S's isn't in a position to behaviourally indicate, either verbally or non-verbally, what spatiotemporal relation the contents c_1 of e_1 and c_2 of e_2 stand in.

C1: Therefore, e_1 and e_2 aren't spatiotemporally unified.

P1 is the claim that S *actually has* an experience both of the question mark (e_1) and the dollar sign (e_2). The claim is supported by the fact that (i) objects are presented both in *S*'s RVF and S's LVF, (ii) *S* can verbally report the presence of the question mark and non-verbally report the presence of the dollar sign (through a left-handed drawing), and (iii) both LH and RH are capable of supporting consciousness. *Prima facie*, at least, the best explanation of (i)-(iii) is that *S* has two experiences.

Even though this motivation of P1 is straightforward and intuitive, there are ways of objecting to the premise.[21] An example of an objection of this kind will be discussed in section 8.2. For now, though, let's assume that P1 is true.

If the notion of unity under consideration is Spatial Unity, the stipulation that e_1 and e_2 are simultaneous, in the sense of having contents c_1 and c_2 that are experientially presented to *S* as being contemporaneous, must be added to P1. For, as mentioned in section 4.5, contents of experience can be experientially presented as spatially related only if the contents are presented as simultaneous.

On the other hand, if the relevant form of unity is Temporal Unity, P1 must be kept free of the assumption that c_1 and c_2 are presented as contemporaneous to *S*. For if they were, then they would rank as temporally unified, and the argument would self-debunk before getting past its first premise. Instead, the relevant addition to P1, if the argument targets Temporal Unity, would be a stipulation that the vehicles of e_1 and e_2 occur in close temporal succession (see section 4.2).

Let's move on to P2—arguably the most important premise of the argument. According to P2, e_1 and e_2's being spatiotemporally unified requires that *S* is able to verbally or non-verbally (e.g., through drawing) *report* the spatiotemporal relation between the question mark and the dollar sign (a right-left relation, and a simultaneity relation, respectively). Another way of expressing the premise is as the principle that, in order for e_1 and e_2 to be spatiotemporally unified, e_1 and e_2 must be *access unified* in the right way.

[21] More specifically, we'll consider Bayne's objection (2010) to there being two experiences here, or at the very least two *simultaneous* experiences.

It should be noted that P2 doesn't need to be construed as the claim that it's metaphysically necessary that if e_1 and e_2 are spatiotemporally unified, they're access unified. Rather, a more attractive interpretation of P2 would be as the claim that Access Unity, of the relevant kind, between e_1 and e_2 is the only, or most significant, evidence there could be of the experiences' being spatiotemporally unified. Consequently, provided that e_1 and e_2 *aren't* access unified in the right way, the best explanation is that e_1 and e_2 aren't spatiotemporally unified either.

P3 is the negation of the consequent of P2. It's the claim that e_1 and e_2 aren't access unified for *S*. This matches the experimental data. *S* is entirely incapable of making comparisons between the question mark and the dollar sign. In no way, shape, or form does her verbal and non-verbal behaviour betray an awareness of a spatiotemporal relation between the question mark and the dollar sign, nor even an awareness of being visually presented with two objects. Thus, it follows from P2 and P3 that e_1 and e_2 aren't spatiotemporally unified (C). Therefore, in certain experimental contexts, SB-patients have experiences that are spatiotemporally disunified.

P2 seems to be the main premise that a defender of Spatiotemporal Unity could object to. One such objection could go as follows: "Neurotypical subjects sometimes have spatiotemporally unified experiences even though the subjects are incapable of specifying the spatiotemporal relation between the experiences' contents. However, their experiences are nevertheless spatiotemporally unified (or so we ordinarily assume). Since neurotypical subjects' experiences can be spatiotemporally unified even though the spatiotemporal relations between their contents are unspecifiable, it follows that the same could be true of SB-patients."

A spatiotemporal relation between experience-contents that isn't describable by a neurotypical subject will presumably have to do be a fine-grained, highly determinate relation. To take a temporal example, suppose that a neurotypical subject listens to an already-fast piece of music at ten times its usual speed. There may be an instance where the subject isn't able to specify, either through verbal or (less plausibly) non-verbal behaviour, whether or not two auditory contents c_1 and c_2 are experientially presented as simultaneous. However, the subject's failed specification attempt hardly constitutes a reason for inferring that *there is no* spatiotemporal relation between c_1 and c_2. So, likewise, one could argue that SB-patients' inability to specify the spatiotemporal relation between the contents of their experiences doesn't demonstrate that the experiences are spatiotemporally disunified.

However, this objection to P2 entirely misconstrues the demand that the Spatiotemporal Unity of two experiences requires that the subject is capable of behaviourally indicating what the spatiotemporal relation between the experiences' contents are. What's at issue here aren't fine-grained, highly determinate spatiotemporal relations. Rather, the spatiotemporal relations that are pertinent for the argument are some of the most basic and coarse-grained one can imagine; in the spatial case, relations such as "is to the left of," and in the temporal case, coarse-grained simultaneity. While it's probably true that the existence of spatiotemporal relations between the contents of experience doesn't require that fine-grained features of these relations are reportable, it *is* a reasonable condition on experience-contents' standing in spatiotemporal relations that the subject is able to behaviourally indicate what the most basic features of these relations are. So the objection fails.

The basic structure of the argument will be the same as the argument for SB-patients' having spatiotemporally disunified experiences:

P1: When a SB-patient S is simultaneously tachistoscopically presented with visual stimuli in RVF and LVF, then S has two experiences e_1 and e_2, with e_1 corresponding to the visual stimuli in RVF and e_2 corresponding with the visual stimuli in LVF.

P2: If e_1 and e_2 are phenomenally unified, then *S* is in a position to behaviourally indicate, either verbally or non-verbally, that she has e_1 *and* e_2.

P3: *S* isn't in a position to behaviourally indicate, either verbally or non-verbally, that she has both e_1 *and* e_2.

C1: Therefore, e_1 and e_2 aren't phenomenally unified.

The main difference between this argument and the argument for in the preceding subsection has to do with P2 (and, resultantly, P3). Here, e_1 and e_2 being unified doesn't require that S is able to behaviourally indicate what specific kind of relation the contents stand in. Rather, the relation between e_1 and e_2 which *S* must be able to express, either linguistically or non-linguistically, is the relation of "occurring together." *S* must, that is, be able to express that she has both e_1 *and* e_2.

This is a weaker requirement than in the corresponding premise of the argument against the Spatiotemporal Unity of SB-patients' consciousness. For we can imagine a

subject who's able to articulate that she has two experiences whose contents are experientially presented together, but who isn't able to specify the spatiotemporal relations between the contents. Such a subject may say something along the following lines:

> There were two things going on—a flash of blue and a flash of red. I have no idea of what came first and what came after, and I don't know "where" the blue flash was and "where" the red flash was. But it definitely felt like something. It was very strange—not at all like what I feel when I just experience a flash of blue, or when I just experience a flash of red.

This example shouldn't be interpreted as being committed to the claim that Phenomenal Unity is possible without Temporal Unity. The example is simply intended to show that there are hypothetical phenomenological descriptions—possibly of impossible experiential situations—that, if uttered by subjects, would count as evidence for Phenomenal Unity, but not for Spatial Unity or Temporal Unity. Therefore, P2 is formulated in a different, more permissive way in the argument under consideration in this subsection than in the argument for SB-patients' having spatiotemporally disunified experiences.

Ultimately, though, this difference between the arguments counts for very little. For the fundamental challenge presented by SB-patients to Spatiotemporal Unity applies to Phenomenal Unity as well: SB-patients' utter incapacity to compare simultaneous, tachistoscopically presented visual stimuli. Just like this is reflected in SB-patients' inability to behaviourally indicate the spatiotemporal relations between the

contents of e_1 and e_2, it's also expressed in the patients' incapacity of expressing or otherwise showing that they have two experiences at all. Thus, P3 is true, and, given P2, it follows that SB-patients have experiences which aren't phenomenally unified (C).

Objection: Phenomenal Unity without Access Unity

P2 articulates a connection between the Phenomenal Unity of e_1 and e_2 and the experiences' being access unified. But why accept that Phenomenal Unity entails Access Unity? Couldn't SB-patients' experiences be phenomenally unified even though they aren't access unified?

It's unclear what claim is being made in this objection. On one reading, the claim is that Access Unity isn't conceptually or metaphysically necessary for Phenomenal Unity. That is true. However, it doesn't show that P2 is false. For P2 can be true without being conceptually or metaphysically true. In order to undermine P2, the objection must show that, in the specific case of SB-patients, there are reasons to think that phenomenally unified experiences can fail to be access unified. And there aren't any such reasons. In general, presence or absence of Access Unity between experiences is the only evidence there can be of presence or absence of Phenomenal Unity.[22]

In summary, SB-patients have spatially disunified experiences, temporally disunified experiences, and phenomenally disunified experiences. Hence, the Unified One-Stream View of SB-patients is false.

[22] If there are general philosophical reasons for thinking that it's impossible for contemporaneous experiences in a stream of consciousness to not be phenomenally unified, then there may be reasons for inferring that experiences are phenomenally unified even though they aren't access unified. We'll touch on considerations like this in section 10, where we'll discuss the transitivity of the unity relation.

8.2 Bayne's Defence of Phenomenal Unity

Bayne's Switching View

In chapter 9 of *The Unity of Consciousness*, Bayne (2010) defends the view that SB-experiments don't disprove the thesis that simultaneous phenomenally conscious experiences of a subject must be phenomenally unified. Hence, he defends a view on which all of SB-patients' simultaneous experiences are phenomenally unified. In Bayne's view is that consciousness is necessarily phenomenally unified, for any subject of experience:

> Necessarily, for any conscious subject of experience (*S*) and any time (*t*), the simultaneous conscious states that *S* has at *t* will be subsumed by a single conscious state—the subject's total conscious state. (Bayne, 2010: 16.)

According to Bayne, attempts of inferring lack of Phenomenal Unity from SB-patients' behaviour in SB-experiments fail because they mistakenly assume that, when a SB-patient is simultaneously tachistoscopically presented with visual stimuli in RVF and LVF, the resultant experiences, associated with LH and RH respectively, are simultaneous. On Bayne's view, such experiences *aren't* simultaneous. A SB-patient's cerebral hemispheres aren't conscious simultaneously but rather sequentially:

> Rather than suppose that the patient's two hemispheres are conscious in parallel, we should think of consciousness in the split-brain as moving or switching from one hemisphere from another. Although both hemispheres can process information concurrently, they take turns supporting consciousness. In effect, the switch model paints patients as suffering from a kind of fluctuating perceptual extinction: when the left (right) hemisphere is activated stimuli in the LVF (RVF) are typically ignored in favour of stimuli in RVF (LVF). The patient might be conscious of [the dollar sign] (due to right hemisphere activity), or she might be conscious of [the question mark] (due to left hemisphere activity), but she will not be conscious of both [the dollar sign] and [the question mark] *at once*, even when the two words are simultaneously presented to her. (Bayne, 2010: 210.)[23]

[23] In the quote, Bayne discusses as SB-experiment that hasn't been introduced here. The substitutions adapt what he says to one of the experiments introduced in section 6.2

On this view, a SB-patient's stream of consciousness "jumps" between her hemispheres during SB-experiments. That is, for SB-patients, LH and RH aren't simultaneously consciousness-supporting but rather successively consciousness supporting. So, for example, when a question mark and a dollar sign are simultaneously tachistoscopically presented to a SB-patient *S*, in RVF and LVF respectively, *S* doesn't simultaneously have an experience e_1 of the question mark and an experience e_2 of the dollar sign. Instead, either *S* doesn't consciously experience one of the stimulus items at all, or *S* perceives both stimulus items but at slightly different points in time.[24,25] In either case, the distinctive feature of Bayne's view is that the stream of consciousness is continuously interhemispherically switching in SB-experiments, and that, as a result, *S* never has two simultaneous but phenomenally disunified experiences.[26]

Bayne writes that the continuous interhemispheric "switching" of *S*'s stream of consciousness explains *S*'s behaviour. When *S* mentions having seen a question mark, LH is supporting *S*'s stream of consciousness, meaning that *S* can verbally report the previously presented LH-associated visual stimuli in RVF. Conversely, when *S* draws a question mark using her left hand, *S*'s stream of consciousness has moved to RH,

[24] If the latter option holds, *S* experiences e_1 at t_1, while LH realizes her stream of consciousness, after which *S* experiences e_2 at t_2, when the realization base of *S*'s stream of consciousness has "moved" to RH.
[25] Bayne isn't very explicit about if he believes that SB-patients experience *both* the objects in RVF and the objects in RVF, at different times, or whether he believes that the patients only experience one of the objects. One consideration speaking in favour of the former option is that tachistoscopic presentation of visual stimuli already is very quick (a tenth of a second) so as to preclude the possibility of eye movements. It's unclear that interhemispheric switching could occur and result in the formation of two experiences within such a short range of time.
[26] In order for Bayne's proposal to make sense, the temporal relations are between experience-vehicles rather than between experience-contents.

meaning that *S* can non-verbally report the previously presented RH-associated visual stimuli in LVF:

> [...] access disunity arguments assume that the patient's two hemispheres must be simultaneously conscious because the stimuli are simultaneously projected to the patient's visual hemifield, and because each hemisphere can respond to the stimulus in its hemifield as and when required. But both inferences are contentious. Perhaps the ability of patients to respond in this way is the result of consciousness switching rapidly and effortlessly between hemispheres in response to the demands of the patient's context. The hemisphere that is silent on any one trial may be so because it is unconscious rather than because it is unable (or unwilling) to 'speak'. (Bayne, 2010: 213, footnote omitted.)

Let's call Bayne's view of SB-patients the *Switching View*. The Switching View entails that SB-patients don't have simultaneous experiences in SB-experiments involving tachistoscopic presentation of visual stimuli. To be sure, pairs of such experiences, e_1 and e_2, aren't phenomenally unified. However, this doesn't imply that there is a lapse in the Phenomenal Unity of SB-patients' consciousness. For Phenomenal Unity is defined as a synchronic relation. Therefore, e_1 and e_2 aren't even eligible *candidates* for being phenomenally unified.

Response to Bayne's Argument

Positing a mechanism for the continuous "switching" of a stream of consciousness between hemispheres only ensures that the vehicles of SB-patients' experiences occur at different points in time. Thus, that mechanism ensures the non-simultaneity of SB-patients' experiences only given a vehicular construal of simultaneity. However, the conception of simultaneity that's most relevant for Phenomenal Unity isn't vehicular simultaneity but content-simultaneity. Interhemispheric switching of the support base

for streams of consciousness doesn't guarantee the non-simultaneity of SB-patients' experiences.

However, as a measure of charitableness towards Bayne, let's assume that the pertinent conception of simultaneity for Phenomenal Unity is vehicular simultaneity. Given this assumption, the Switching View offers a cogent explanation of SB-patients' behaviour which is consistent with Bayne's view that contemporaneous experiences of a subject must be phenomenally unified.

Even though the Switching View is cogent, it's inferior to the view of SB-patients defended in this thesis, the *Malleable Stream View*. The Switching View is inferior because it's more complex than the Malleable Stream View, and not more explanatory. Both when it comes to explaining the behaviour of SB-patients intra-experimentally and explaining their behaviour intra-experimentally, the Switching View adds superfluous complexity.

Let's begin by looking at the intra-experimental case with a SB-patient *S* having tachistoscopically induced experiences e_1 and e_2. Both the Switching View and the Malleable Stream View will describe e_1 and e_2 as not being access unified. The Malleable Stream view will explain the absence of a synchronic Access Unity relation between e_1 and e_2 as resulting from *S*'s lack of an intra-cranial mechanism for interhemispheric information integration, in conjunction with the unavailability of opportunities for behavioural cross-cuing. The Switching View explains the failure of Access Unity between e_1 and e_2 in the same way.[27] The only thing the Switching View adds to the Malleable Stream View is an interhemispheric switching mechanism. That mechanism is

[27] This would be a diachronic Access Unity relation.

only there to ensure the non-simultaneity of e_1 and e_2, and plays no distinctive role in explaining *S*'s behaviour.

Let's move on to the case of *S*'s extra-experimental behaviour. As mentioned in section 5, SB-patients generally behave comparatively normally outside of the artificial constraints of SB-experiments. Bayne discusses that this may be because, outside of SB-experiments, there are rarely switches in what hemisphere is supporting consciousness, with LH appropriating the stream of consciousness. As an alternative possibility, he also mentions that interhemispheric switches in the "seat" of consciousness could be so rapid that it produces the impression that LH and RH are consciousness-supporting simultaneously.

However, regardless of which of these options is more promising, information integration between the hemispheres will only be possible through behavioural cross-cuing. For a dominant LH or successively consciousness-supporting LH and RH don't change the fact that, outside the realm of certain kinds of sensory information that can be interhemispherically integrated subcortically, cross-cuing is the only way for SB-patients to transmit sensory information across the hemispheres. So the Switching View will have to appeal to behavioural cross-cuing such as eye movements to account for SB-patients' behaviour extra-experientially, just as the Malleable Stream View. Behavioural cross-cuing does all the work in accounting for the normality of SB-patients' behaviour extra-experimentally.

To summarize, Bayne's postulation of a switch mechanism determining which hemisphere is consciousness-supporting at a time doesn't help explaining the behaviour of SB-patients. Rather, it appears to be posited just to ensure the non-simultaneity of

SB-patients' tachistoscopically induced experiences and, by extension, to safeguard the thesis that simultaneous experiences necessarily are phenomenally unified. The Malleable Stream View is simpler, and hence preferable.

CHAPTER 4: MULTIPLICITY AFTER ALL?

9 Multiplicity Through Lack of Unity?

9.1 Unity NSP and Multiplicity

If the arguments in Part 3 are correct, SB-patients have experiences that aren't unified either phenomenally, spatially, or temporally (nor, for that matter, in terms of access). At this stage of the argument, that doesn't tell us how many streams of consciousness SB-patients have. For if experiences don't satisfy the antecedent of Unity SSPs, that only establishes that they don't satisfy a sufficient condition for singularity. But that's consistent both with SB-patients' having one disunified stream of consciousness and with SB-patients' having multiple streams of consciousness.

However, the situation is different if the disunities of SB-patients' consciousness are interpreted in the light of *Unity NSPs*. Recall from section 2 that Unity NSPs are principles stating that a set of experiences $e_1, \ldots, e_n$ belong to a stream of consciousness E at t only if $e_1, \ldots, e_n$ are unified (in some respect) at t. That is, Unity NSPs construe unity as a necessary condition for singularity. Consequently, if the arguments in previous sections are correct, and SB-patients have simultaneous experiences which aren't phenomenally, spatially, or temporally unified, then an appropriate form of Unity NSP entails that not all of SB-patients' experiences belong to a single stream of consciousness. Given the additional assumption that any experience belongs to a stream of consciousness, it follows that SB-patients have at least two streams of consciousness.

Subscribing to both Unity SSP and Unity NSP makes unity the mark of co-membership in a stream of consciousness.[28] On this combination of views, unity *defines* streams of consciousness.

9.2 Schechter's Argument for Multiplicity

A commitment to Unity NSP, or something like it, is needed for the success of arguments for Two-Stream Views of SB-patients. To show this, we'll discuss Elizabeth's Schechter's (2018: ch. 2) argument for there being two streams of consciousness (and, in fact, two subjects of experience) corresponding to SB-patients. Schechter's is a good example of the Two-Stream View, as she's unusually explicit about the role that assumptions about unity play in establishing her view.

Schechter's argues for her view of SB-patients in two steps. In the first step (Schechter, 2018: 25-46), she argues that SB-patients have two streams of consciousness (what she calls "experiential perspectives"; 2018: 25.) In the second step of the argument (Schechter, 2018: 46-49), she infers that SB-patients are composed of two subjects of experience, one corresponding to each stream of consciousness. There is only room to discuss the first step of the argument here.

It should be noted that, according to Schechter, "RH elements of experience and LH experiences are systematically unified intrahemispherically and disunified interhemispherically, *both at a time and across time*." (Schechter, 2018: 28.) That is, on

[28] It will be assumed here that views combining Unity SSP and Unity NSP define those principles in terms of the *same* kind of unity relation. However, it isn't mandatory to do this. Conceivably, one could subscribe to Unity SSP for one form of unity and subscribe to Unity NSP for another form of unity, without the two forms of unity being the same.

Schechter's view, SB-patients' consciousness is interhemispherically disunified and intrahemispherically disunified both synchronically and diachronically. However, Schechter's argument touching on diachronic conscious unity won't be discussed here.

Schechter claims that the following ingredients are needed for a persuasive argument for the Two-Stream View (or a "2-perspectives claim, as she calls it"):

> Making the case for the 2-perspectives claim requires showing three things. First, we need to establish that both hemispheres are the source or site of elements of experience. I call this the *bihemispheric consciousness claim*. Second, we need to determine that RH elements of experience are unified with each other, and that LH elements of experience are unified with each other. This *intrahemispheric unity claim* means that a split-brain subject has *at most* two perspectives. The *interhemispheric disunity claim* then shows that a split-brain subject has *exactly* two perspectives: RH elements of experience systematically fail to be unified with LH elements and thus they are not elements of a common perspective. (Schechter, 2018: 28-29.)

So, according to Schechter, an argument for the Two-Stream View must establish the truth of the bihemispheric consciousness claim, the intrahemispheric unity claim, and the interhemispheric disunity claim. For our purposes, it'll be helpful to go through the claims in reverse order. (A quick note about the bihemispheric consciousness claim: It actually requires the hemispheres to be *simultaneously* supporting consciousness. More will be said on this below.)

Schechter's *interhemispheric disunity claim* is that, for SB-patients, RH-associated experiences systematically fail to be access unified with LH-associated experiences. Schechter argues for this position in roughly the same way as I do in section 8. SB-patients' behaviour shows that their experiences aren't interhemispherically access unified, and the absence of interhemispheric Access Unity can be used to infer the absence of other forms of unity relations such as Phenomenal Unity. On this basis, Schechter derives that not all of SB-patients' (simultaneous) experiences belong in a single stream

of consciousness. For this derivation, she relies on the following analogue of Unity NSP:

> There are a few different ways to use the access unity relation to determine how many perspectives a split-brain subject has. The least controversial of these is via a kind of *exclusion principle*, according to which two elements of experience that are not access unified are not elements of a common perspective. (Schechter, 2018: 27.)

The principle Schechter mentions here follows from an access version of Unity NSP. Thus, without an appropriate version of (something like) Unity NSP, Schechter's case for the interhemispheric disunity claims fails.

Let's move on to the *intrahemispheric unity claim*. This is claim consisting of two sub-claims. The first is that, when it comes to experiences that are associated with the *same* hemisphere, SB-patients' experiences are unified. The evidence Schechter cites in support of this claim mainly pertains to each hemisphere's capacity to integrate disparate kinds of sensory information and thereby facilitate the execution of complex behavioural tasks. In turn, the second sub-claim of the intrahemispheric unity claim is that, as a result of their unity, RH-associated experiences belong to a stream of consciousness E_1 while LH-associated experiences belong to a stream of consciousness E_2. Given the interhemispheric disunity claim, it follows that E_1 and E_2 aren't identical.

Unlike above with the interhemispheric disunity claim, Schechter doesn't explicitly define the principle that licenses categorizing RH-associated experiences and LH-associated experiences as belonging to the same (respective) stream of consciousness. However, it's clear that she's relying on an access version of something like Unity SSP.

Let's conclude the discussion of Schechter's argument for the Two-Stream View by examining the *bihemisperic consciousness claim*. Schechter characterizes this as the claim

that "[…] both hemispheres are the source or site of elements of experience." (Schechter, 2018: 28-29.) Interestingly, she later writes that this claim is consistent with Bayne's Switching View of SB-patients, according to which RH and LH never realize consciousness simultaneously, but only successively (Schechter, 2018: 35.)

However, the bihemispheric consciousness claim is better construed as requiring that SB-patients *simultaneously* have RH-associated experiences and LH-associated experiences. For if the claim is construed in this way, then it could serve as a dividing line between Schechter and Bayne, helping to explain why they endorse views of SB-patients that are so different. Just like Schechter, Bayne accepts Unity NPS. On Bayne's view, it's *impossible* for there to be simultaneous phenomenally conscious experiences of a subject that aren't phenomenally unified (Bayne, 2010: 16). Therefore, the only option Bayne has for avoiding having to adopt a Two-Stream View is to repudiate the (updated) bihemispheric consciousness claim, thereby denying the simultaneity of interhemispherically disunified experiences. Without this move, and the concomitant postulation of a switching mechanism for the hemispheric support base of consciousness, Bayne's view would collapse into Schechter's.

As a summary of Schechter's argument for the Two-Stream View of SB-patients, she utilizes Unity NSP to infer that SB-patients don't have a single stream of consciousness, after which she employs Unity SSP to infer that SB-patients have two streams of consciousness. Together, Unity NSP and Unity SSP create a route from interhemispheric disunity to a multiplicity of streams of consciousness.

9.3 Nontransitivity as a Counterexample to Unity NSP

Schechter's argument for the Two-Stream View doesn't work without a commitment to Unity NSP. Hence, if a relevant counterexample to Unity NSP can be identified, Schechter's argument would fail.

A counterexample to Unity NSP would be an instance of two (simultaneous) experiences e_1 and e_2 that aren't unified but which nevertheless belong to the same stream of consciousness. An extreme case of this would be a stream of consciousness in which none of the experiences were unified with each other. However, it's not obvious that such streams are possible. And in any case, SB-patients don't have such utterly disunified streams of consciousness. SB-experiments patients have experiences exhibiting a significant degree of unity, intrahemisperically in SB-experiments, and both intrahemispherically and interhemispherically outside of SB-experiments. Hence, we need another counterexample to Unity NSP, one that better reflects what one conceivably could expect to find out about SB-patients' consciousness.

In order to identify a relevant counterexample to Unity NSP, it'll be instructive to consider the formal properties of the unity relation. If Unity NSP is true, unity is an equivalence relation that's symmetric, reflexive, and transitive. Unity is clearly both symmetric and reflexive.[29] That leaves us with *transitivity*: The condition that if e_1 is unified with e_2, and e_2 is unified with e_3, then e_1 is unified with e_3. If the unity relation were *nontransitive*, then Unity NPS would be false. That, in turn, would mean that

[29] Regarding the reflexivity point, the following remarks can be made. Firstly, any spatial content of an experience is experientially presented as being spatially co-located with itself (Spatial Unity). Secondly, any content of an experience is experientially presented as being simultaneous with itself (Temporal Unity). Thirdly, any content of an experience is experienced by a subject together with itself (Phenomenal Unity). Fourthly, any experience-vehicle that's available for rational control of action, speech, and inference is available together with itself (Access Unity).

Schechter's interhemispheric disunity claim lacks justification. In other words, SB-patients' having simultaneous experiences that aren't unified wouldn't show that they don't have a single stream of consciousness.

A nontransitive unity relation would be a counterexample to Unity NPS. And in this case, it would be a relevant counterexample. For a reasonable example of nontransitivity isn't one involving no unified experiences, but rather one in which *some but not all* experiences are unified.

Thus, a lot hinges on whether unity is transitive or non-transitive. Are there any persuasive grounds for holding that unity *must* be transitive?

10 Must Unity be Transitive?

In this section, I'll argue that unity is a nontransitive relation. The argument proceeds indirectly, through a discussion of arguments for the transitivity of the unity relation.[30] I'll show that these arguments fail. The failure of the arguments is then inferred to constitute evidence for the nontransitivity of the unity relation.

Let's call the thesis that unity is transitive "Transitivity" (with a capital "T"). Correspondingly, let's denote the thesis that unity is nontransitive with "Nontransitivity" (with a capital "N"):

[30] For more direct defences of the intransitivity of the unity relation, see (Lockwood, 1989: ch. 6; Tye, 2003: ch. 5).

Transitivity: Necessarily, for any experiences e_1, e_2, and e_3, and any point in time t, if e_1 is unified with e_2 at t, and e_2 is unified with e_3 at t, then e_1 is unified with e_3 at t.

Nontransitivity: It's possible that, for some experiences e_1, e_2, and e_3, and some point in time t, e_1 is unified with e_2 at t, and e_2 is unified with e_3 at t, but e_1 is not unified with e_3 at t.

10.1 Unimaginability and Inconceivability

A counterexample to Transitivity would be if a subject S had simultaneous experiences e_1, e_2, and e_3, where e_1 is unified with e_2, and e_2 is unified with e_3, but e_1 isn't unified with e_3. Such counterexamples are difficult to imagine. For example, focusing on Phenomenal Unity, it's difficult to see how *S* could, at the same point in time, experience contents c_1 and c_2, and experience contents c_2 and c_3, without thereby experiencing c_1, c_2, *and* c_3 as occurring together. It may be thought that the (alleged) unimaginability of counterexamples to Transitivity shows that instances of intransitivity are impossible, or at least that there aren't any good reasons to think that Intransitivity is possible.

Both Schechter and Bayne present versions of these arguments (Bayne, 2010: 36-45; Schechter, 2018: 46-49). Admittedly, Schechter doesn't argue for the unimaginability of experiences' violating Transitivity but rather for the unimaginability of a single subject of experience having two streams of consciousness.[31] However, the reason for the unimaginability will be the two in the same cases, namely our inability to imagine

[31] On this basis, Schechter infers that SB-patients must be composed of multiple subjects of experience.

what it would be like to be a Transitivity-violating or multi-stream subject, an incapacity to "[…] project ourselves into their [the subjects'] mental lives." (Nagel, 1971: 410.)[32]

In other words, what's allegedly unimaginable about Intransitivity is the experiential situation of the subject whose experiences violate Transitivity. That is, the unimaginability at issue is what we might call first-personal, phenomenological unimaginability. It's impossible, at least on the face of it, to imagine what it would be like to be a subject whose experiences don't satisfy Transitivity. As Bayne writes "[we] are unable to project ourselves into a partially unified phenomenal perspective." (Bayne, 2010: 39.) A simple unimaginability argument against Nontransitivity would have the following form:

P1: A subject whose experiences don't conform with Transitivity is unimaginable.

P2: Whatever is unimaginable is impossible.

C: Therefore, a subject whose experiences aren't transitively unified is impossible.

There are two ways to disambiguate P2, and on neither disambiguation is the argument sound. On the first disambiguation, P2 is the claim that whatever is unimaginable *for human beings* is impossible. That claim is clearly incorrect. There are a range of creatures whose consciousness human beings aren't in a position to positively imagine. Famously,

[32] Schechter also cites this section from Nagel (2018: 49).

it's often thought that we cannot know what it's like to be a bat (Nagel 1974). It doesn't follow that bats are impossible creatures.

In turn, on the second disambiguation, P2 is the claim that whatever is unimaginable *simpliciter*, for any possible being, is impossible. This version of the claim is more plausible than the first version, although it's far from uncontroversial. But on this understanding of "unimaginability," P1 isn't adequately justified. For the most we're entitled to claim here is that subjects whose experiences violate Transitivity are unimaginable for subjects whose experiences (presumably) conform with Transitivity. That doesn't entail unimaginability simpliciter.

So an imaginability argument against Nontransitivity must have a more subtle form than the argument above. Bayne is aware of this, and doesn't argue against Nontransitivity in the way above.[33] Instead, he uses considerations about unimaginability to construct an *inconceivability argument* against Nontransitivity. Borrowing two terms from (Van Cleve, 1983), Bayne distinguishes between *strong inconceivability* and *weak inconceivability*. He writes that "[a] scenario is strongly inconceivable for *S* when *S* seems to see that it is impossible, whereas a scenario is weakly conceivable for *S* when *S* cannot see that it is possible." (Bayne, 2010: 43.) Bayne then goes on to argue that Nontransitivity is weakly inconceivable. Furthermore, he argues that the weak inconceivability of Nontransitivity means that there's cause for significant scepticism about the possibility of counterexamples to Transitivity:

> Is [Nontransitivity] *strongly* inconceivable? That seems unlikely. Transitivity may be a deep feature of synchronic unity, but I very much doubt that it reveals itself to us as such [---]. That leaves weak inconceivability. Is [Nontransitivity] weakly inconceivable? I am

[33] Schechter probably doesn't endorse a straightforward unimaginability-impossibility link either. However, the logical structure of her argument is somewhat opaque.

> inclined to think that it is. At any rate, I cannot 'see' that [Nontransitivity] is possible, and as far as I can tell I am not alone in this [---]. The prudent position, it seems to me, is to retain partial unity as a potential model of consciousness, albeit one that is surrounded by a significant degree of suspicion. (Bayne, 2010: 43-44.)

Bayne's position here is that, while we lack decisive reasons for the impossibility of Nontransitivity, we nevertheless cannot conceive of what it would take for Nontransitivity to be true. He concludes by writing that while Nontransitivity shouldn't be ruled out entirely, it should nevertheless be regarded with "a significant degree of suspicion."

There are two ways to question this argument. Firstly, one could question whether Nontransitivity really is weakly inconceivable. Secondly, one could question Bayne's contention that Nontransitivity should be treated with suspicion just on accounts of being weakly inconceivable. As we'll see, both these strategies are promising.

10.2 The Weak Inconceivability of Nontransitivity

Bayne doesn't provide much defence of Nontransitivity's being weakly inconceivable, besides citing the autobiographical and sociological fact that he "[…] cannot 'see' that [Nontransitivity] is possible, and […] [isn't] alone in this." (Bayne, 2010: 44.) It's up to us to fill in the blanks here. It's possible that what Bayne has in mind is connected with the unimaginability of subjects whose experiences violate Transitivity. Unlike in the argument above, Bayne would rely on the unimaginability of such subjects to motivate their weak inconceivability, rather than their impossibility.

In the absence of specificity from Bayne, this is how we'll understand his argument. The question now becomes whether, in the case of Transitivity-violating subjects of experience, unimaginability entails weak inconceivability.

It doesn't. To see this, let's consider the following hypothetical statement made by a subject *S* converted to the view of SB-patients defended in this thesis, namely the Malleable Stream View:

> The Malleable Stream View is the theory that most conservatively and non-dogmatically accounts for the behaviour of SB-patients both inside and outside of experimental settings. I can thus see how the theory is true—indeed, I have a strong conviction that it *must* be true. Since the theory entails that unity cannot be transitive, I have a correspondingly strong conviction that the unity relation is nontransitive. Admittedly, it's difficult for me to imaginatively project myself into the position of someone whose experiences don't satisfy the condition of Transitivity. However, given the attractiveness of the Malleable Stream View, this simply shows that there must be something confused about how I think of consciousness.

Counterexamples to Transitivity are unimaginable for *S*. Note that, here, "imaginability" has to be understood in a distinctively first-personal, phenomenological sense. The unimaginability of counterexamples to Transitivity, for a subject, consists in the subject's incapacity of projecting herself into the experiential perspective of subjects whose experiences violate Transitivity.

Now, given that Nontransitivity is *unimaginable* for *S* in this, first-personal, phenomenological sense, does it follow that Nontransitivity also is *weakly inconceivable* for *S*? The answer is *no*. For *S* fulfils any reasonable criteria for "seeing" how Nontransitivity

could be true. For unlike imagining a subject whose experiences violate Transitivity, conceiving of such a subject doesn't require the ability of projecting oneself into the subject's experiential perspective; it doesn't require knowing what it's like to be the subject. *S* recognizes how rejecting Transitivity fits into the theory that constitutes the best explanation (let's suppose) of an explanatory domain, namely that of SB-patients. Therefore, *S* can conceive of there being counterexamples to Transitivity. In other words, unimaginability doesn't entail weak inconceivability, and, equivalently, conceivability doesn't entail imaginability.[34]

The only way to avoid this conclusion would be to insist that the only form of conceivability that matters when it comes to alleged counterexamples to Transitivity is imaginability. In effect, this would involve an identification of conceivability and imaginability. But such an identification would be misguided. If *S* has access to convincing, broadly third-personal evidence of an abductive nature indicating that there must be counterexamples to Transitivity, then there's a perfectly good sense in which Nontransitivity is conceivable for *S*.

In conclusion, the unimaginability of Nontransitivity doesn't ensure its weak inconceivability. Therefore, Bayne's claim that Nontransitivity is weakly inconceivable is inadequately justified.

[34] I'm assuming that the negation of weak inconceivability is conceivability. However, there might be a distinction between strong conceivability and weak conceivability. Depending on how these notions are characterized, the situation *vis-à-vis* entailment relations would be more complicated than how it's presented to be here.

10.3 Weak Inconceivability and the Unity of Imaginative Awareness

Suppose (implausibly) that Nontransitivity is conceivable only if it's possible to imagine what it's like being a subject with experiences violating Transitivity. In that case, Bayne would be right about Nontransitivity being weakly inconceivable. But why assume that this establishes that views rejecting Transitivity should be "[...] surrounded by a significant degree of suspicion?"(Bayne, 2010: 44.)

In this section, I'll present an argument for the evidential insignificance of Nontransitivity's being weakly inconceivable. That is, I'll argued that, even if Nontransitivity is weakly inconceivable, that has no bearing on whether it should be endorsed or not. Here's the argument:

P1: The only way for a subject *S* to conceive what it's like to be a subject *S** whose experiences e_1, e_2, and e_3 violate Transitivity is by imagining having e_1, e_2, and e_3 simultaneously.

P2: Imagining a set of simultaneous experiences imposes unity on those experiences.

C: Therefore, it's impossible for *S* to conceive what it's like to be S*.

It's assumed here that imaginability is the only relevant form of conceivability, at least when it comes to evaluating claims about Transitivity and Nontransitivity. This was questioned in the last section, but here we're granting the premise for the sake of charitableness. Provided that imaginability is the only relevant form of conceivability, then

the first premise of the argument is true: The conceivability of a subject whose experiences violate Transitivity requires being able to imagine what it's like being that subject.

P2 is the interesting premise of the argument. According to P2, imagination has a *unity-conferring* character. We can motivate this premise by making a comparison with the mechanism of introspection.

Introspection goes hand-in-hand with unity. If a subject is in a position to make introspective comparisons between two simultaneous experiences e_1 and e_2, then e_1 and e_2 must be unified. One explanation of the connection between introspection and unity is that a subject is able to make introspective comparisons between e_1 and e_2 in virtue of their unity. However, on an alternative view, there are cases where the unity of e_1 and e_2 is grounded in the subject introspecting them, or in her capacity to introspect them. In this sense, it may be thought that introspection *impose* unity on experience.

It's not unbelievable that imagination could have a similar character. One could hold that, when one imagines two experiences e_1 and e_2, the faculty of imagination imposes unity on e_1 and e_2. If this were true, e_1 and e_2 would be unified in virtue of the unity-imposing character of imagination. Imagining e_1 and e_2 would involve entering a state of imaginative awareness whose content encompasses the contents of e_1 and e_2.

Suppose that P2 is a true principle about imagination. It then follows that it's *trivially impossible* to imagine what it's like being a subject whose experiences don't conform with Transitivity. In order to imagine what it's like to be such a subject (S^*), S must imagine simultaneously undergoing e_1, e_2, and e_3, with e_1 being unified with e_2 and e_2 being unified with e_3, but without e_1 and e_3 being unified. However, if imagination automatically imposes unity on simultaneous experiences, then e_1, e_2, and e_3 will

all be unified for *S*. There's simply no way for *S* to imagine simultaneously undergoing e_1, e_2, and e_3 without all the experiences standing in unity relations to each other.

Given P1, it follows that it's impossible for *S* to conceive what it's like to be *S** (C). But, given P2, this doesn't have any bearing on the structure of *S**'s consciousness or on what her first-personal, phenomenological perspective is like. Thus, for *S**, e_1 may very well be unified with e_2, and e_2 with e_3, even though e_1 isn't unified with e_3. The impossibility for *S* of conceiving of what it's like to be *S** doesn't tell one anything about whether subjects such as *S**, whose experiences violate Transitivity, are impossible or non-actual. If it's trivially impossible to conceive of the consciousness of a subject whose experiences violate Transitivity, Bayne is wrong in thinking that views positing Nontransitivity should be treated with suspicion.

Both moves in Bayne's unimaginability argument fail. There aren't good reasons to believe that Nontransitivity is weakly inconceivable, and even if there are, it counts for nothing evidentially. In conclusion, if one wants to show that Nontransitivity is impossible, considerations of unimaginability get one nowhere.

10.4 The Uniqueness of What-It's-Likeness

In this section, I'll consider a more direct argument, not for the unimaginability or inconceivability of Nontransitivity, but for its impossibility or incoherence.

The argument starts with the notion of phenomenal consciousness. If a subject *S* is phenomenally conscious, then there's something it's like to be *S*; *S* has a point of view on the world. That is uncontentious. However, one could make a further claim, which may not even seem distinct from the first. This is the claim that if *S* is phenomenally

conscious, there's some *one* thing that it's like to be *S*, some unique point of view that *S* has of the world. Let's call this the *Uniqueness Principle*.

The next step of the argument is to show that there is no unique way that it's like to be a subject *S* whose experiences violate Transitivity. For such a subject, there would be something that it's like to have e_1 and e_2 together, and there would be something that it's like to have e_2 and e_3 together, but there wouldn't be something that it's like to have all of e_1, e_2, and e_3 together. In this sense, there would be no *one* thing that it's like to be *S*. There's wouldn't be a *single* something that it's like to be *S*; rather, *S*' consciousness would be better described by saying that there are some *things* that are like something for *S*.

If this is right, the Uniqueness Principle entails that subjects whose experiences don't conform with Transitivity aren't phenomenally conscious. Hence, if SB-patients are phenomenally conscious, which surely is true, then their experiences cannot violate Transitivity.

The only important premise in this argument is the Uniqueness Principle. The Uniqueness Principle articulates a certain conception of what it is for there to be something that it's like for one to exist. It's not an incoherent conception. However, the important question to ask is whether there are any compelling reasons to accept it. Why suppose that there always must be *one single thing* that it's like to be a phenomenally conscious subject?

An answer would be available if the Uniqueness Principle somehow were part of the ordinary way in which we conceive of consciousness and experiences. However, that's unlikely. Folk-conceptions of mental states are highly indeterminate, and it's

improbable that they involve commitments to such recherché notions as the uniqueness of what-it's-likeness or the transitivity of the unity relation.

Another potential answer would be available if the truth of the Uniqueness Claim were somehow introspectively obvious to one, just in having experience. But that's clearly wrong. Introspective reflection on our experience may reveal that there is one single thing that it's like to be one, one single perspective one has on the world. But that doesn't imply anything about the consciousness of other creatures, including human beings such as SB-patients whose consciousness may have a different structure than that of neurotypical subjects.

So there's no reason to accept that being phenomenally conscious requires that there must be one *single* thing that it's like for one, one single perspective on the world. Phenomenal consciousness simply requires that there's something it's like for one to exist, where "something," has a determinable cardinality and subsumes both single things (such as what is found for neurotypical subjects) and multiple things (such as for SB-patients when undergoing experiences).

The arguments against Nontransitivity considered in this section have all failed. Therefore, I infer that Nontransitivity is defensible. The unity relation may very be nontransitive. This entails the falsity of Unity NSP, and blocks Schechter's argument for the Two-Stream View. Thus, the Malleable Stream View is alive and well.

EPILOGUE

I've argued in this thesis that the Malleable Stream View is the correct theory of SB-patients, and that there aren't any good reasons to endorse either the Unified One-Stream View or the Two-Stream View. In this epilogue, let's briefly take a step back from the particularities of these theories and ask a more general question: What, exactly, should the *desideratum* be for a philosophical interpretation of SB-patients? Should such a theory seek to accommodate previously held theoretical convictions about consciousness? Or should it rather attempt to offer the most conservative and least theoretically demanding explanation of SB-patients' extra-experimental and intra-experimental behaviour, whatever that might mean for one's previously endorsed principles about experience?

Of course, in practise, these two theoretical *desiderata* cannot be held entirely separate. Nevertheless, the distinction between the two *desiderata* captures an important difference between the Unified One-Stream View and the Two-Stream View on the one hand, and the Malleable Stream View on the other hand. For both the former views seem motivated primarily by the first *desideratum* rather than by the second. Both Bayne and Schechter end up endorsing their views because they have theoretical commitments that they're unwilling to give up. In Bayne's case, a commitment to the principle that any experiences of a subject at a time must be phenomenally unified obliges him to posit an interhemispheric switching mechanism to ensure that SB-patients cannot have any simultaneous experiences that aren't phenomenally unified.

Correspondingly, in Schechter's case, Unity NSP, and the fact of interhemispheric disunity between SB-patients' experiences, results in her postulation of two streams of consciousness (and two subjects of experience).

In contrast to the Unified One-Stream View and the Two-Stream View, the Malleable Stream View appears to be primarily motivated by the second theoretical *desideratum* mentioned above. The Malleable Stream View isn't primarily motivated by its conformance with antecedently endorsed principles about the unity of consciousness. Instead, it's primary purpose is accounting for the actual behaviour exhibited by SB-patients in a maximally conservative and parsimonious way. And it does so in a very simple way, by invoking the role that behavioural cross-cuing plays in unifying SB-patients' streams of consciousness outside the artificially controlled parameters of neuropsychological experiments.

I believe that the Malleable Stream View prioritizes the right *desideratum*. The phenomenon of SB-patients is one of the few cases where traditional philosophical questions about the unity of consciousness come in contact with empirical reality. It would be a shame if one's philosophical interpretation of such patients were too rigidly constrained by one's philosophical starting point. I'll conclude this thesis with the following quote from Paul Snowdon, which nicely illustrates what I take to be the fundamental philosophical lesson of the split brain:

> Suppose a psychologist said he was going to run experiments to determine to what degree the experiences of a single human being are linkable and comparable by that human being. Assume too that he suggested that he was doing some experiments involving simultaneous experiences which indicated that […] they could simply not be related by the subject. Would we really say 'That is strange, because we certainly know, without experiment, that some degree of comparability will always be present'? Or would we say 'That is interesting. We are very ignorant about these matters, aren't we?' It seems to me that the second attitude would be our attitude. (Snowdon, 2014: 178.)

BIBLIOGRAPHY

Akelaitis, A. (1941). Psychobiological Studies Following Section of the Corpus Callosum: A Preliminary Report. *American Journal of Psychiatry*, *97*, 1147–1157.

Akelaitis, A. (1944). A Study of Gnosis, Praxis and Language Following Section of the Corpus Callosum and Anterior Commissure. *Journal of Neurosurgery*, *Vol. 1*(3), 94–102.

Alter, T. (2010). A Defense of the Necessary Unity of Phenomenal Consciousness. *Pacific Philosophical Quarterly*, *91*(1), 19–37.

Bayne, T. (2008). The Unity of Consciousness and the Split-Brain Syndrome. *Journal of Philosophy*, *105*(6), 277–300.

Bayne, T. (2010). *The Unity of Consciousness*. New York: Oxford University Press.

Bayne, T. J., & Chalmers, D. J. (2003). What Is the Unity of Consciousness? In A. Cleeremans (Ed.), *The Unity of Consciousness*. Oxford: Oxford University Press.

Block, N. (1995). On a Confusion about a Function of Consciousness. *Brain and Behavioral Sciences*, *18*(2), 227-247.

Bogen, J. E., & Vogel, P. J. (1962). Cerebral Commissurotomy in Man. *Bulletin of the Los Angeles Neurological Society*, *27*, 169–172.

Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.

Davis, L. H. (1997). Cerebral Hemispheres. *Philosophical Studies*, *87*(2), 207–222.

Gazzaniga, M. S. (2005). Forty-Five Years of Split-Brain Research and Still Going Strong. *Nature Review, Neuroscience*, *5*, 653–659.

Gazzaniga, M. S., Bogen, J. E., & Sperry, R. W. (1965). Observations on Visual Perception After Disconnexion of the Cerebral Hemispheres in Man. *Brain*, *Vol 88*(2), 221–236.

Gazzaniga, M. S., & Hillyard, S. A. (1971). Language and Speech Capacity of the Right Hemisphere. *Neuropsychologia*, *Vol. 9*, 273–280.

Gazzaniga, M. S., & Sperry, R. W. (1967). Language after Section of the Cerebral Commissures. *Brain*, *90*, 131–148.

Hurley, S. L. (1994). Unity and Objectivity. In C. Peacocke (Ed.), *Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind* (pp. 49-77). Oxford: Oxford University Press.

Hurley, S. L. (1998). *Consciousness in Action*. Cambridge, Massachusetts: Harvard University Press.

Hurley, S. L. (2003). Action, the Unity of consciousness, and Vehicle externalism. In A. Cleeremans (Ed.), *The Unity of Consciousness* (pp. 78-91). Oxford: Oxford University Press.

Kingstone, A., & Gazzaniga, M. S. (1995). Subcortical Transfer of Higher Order Information: More Illusory than Real? *Neuropsychology*, *Vol. 9*(No. 3), 321–328.

Lockwood, M. (1989). *Mind, Brain, and the Quantum: The Compound 'I'*. Oxford: Blackwell.

Marks, C. E. (1980). *Commissurotomy, Consciousness, and Unity of Mind*. Cambridge, Massachusetts: MIT Press.

McCulloch, W. (1949). Mechanisms for the Spread of Epileptic Activation of the Brain. *Electroencephalography and Clinical Neurophysiology*, *1*, 19–24.

Nagel, T. (1971). Brain Bisection and the Unity of Consciousness. *Synthese*, *22*(3–4), 396–413.

Nagel, T. (1974). What Is it Like to Be a Bat? *Philosophical Review*, *83*(October).

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Peacocke, C. (1983). *Sense and Content: Experience, Thought and Their Relations*. Oxford: Oxford University Press.

Puccetti, R. (1981). The Case for Mental Duality: Evidence From Split-Brain Data and other Considerations. *Behavioral and Brain Sciences*, *4*(1), 93–123.

Schechter, E. (2018). *Self-Consciousness and Split Brains: The Minds' I*. Oxford: Oxford University Press.

Seymour, S. E., Reuter-Lorenz, P. A., & Gazzaniga, M. S. (1994). The Disconnection Syndrome: Basic Findings Reaffirmed. *Brain*, *117*, 105–115.

Shallice, T. (1997). Modularity and Consciousness. In N. Block, O. Flanagan, & G. Güzeldere (Eds.), *The Nature of Consciousness* (pp. 255–267). Cambridge, Massachusetts: MIT Press.

Snowdon, P. F. (2014). *Persons, Animals, Ourselves*. Oxford: Oxford University Press.

Sperry, R. W. (1968). Hemisphere Deconnection and Unity in Conscious Awareness. *American Psychologist*, *23*, 723–733.

Sperry, R. W. (1974). Lateral Specialization in the Surgically Separated Hemispheres. In F. O. Schmit & F. G. Worden (Eds.), *The Neuroscience Third Study Program* (pp. 5–19). Cambridge, Massachusetts: MIT Press.

Sperry, R. W. (1984). Consciousness, Personal Identity and the Divided Brain. *Neuropsychologia*, *Vol. 22*(6), 661–673.

Sperry, R. W., Gazzaniga, M. S., & Bogen, J. E. (1969). Interhemispheric Relationships: The Neocortical Commissures; Syndromes of Hemisphere Deconnection. In P. J. Vinken & G. W. Bruyn (Eds.), *Handbook of Clinical Neurology* (pp. 273–290). Amsterdam: North Holland Publishing Company.

Sperry, R. W., Zaidel, E., & Zaidel, D. (1979). Self-Recognition and Social Awareness in the Deconnected Minor Hemisphere. *Neuropsychologia*, *Vol. 17*, 153–166.

Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.

Tye, M. (2003). *Consciousness and Persons: Unity and Identity*. Cambridge, Massachusetts: MIT Press.

Van Cleve, J. (1983). Conceivability and the Cartesian Argument for Dualism. *Pacific Philosophical Quarterly*, *64*(January), 35–45.

van Wagenen, W., & Herren, R. (1940). Surgical Division of Commissural Pathways in the Corpus Callosum: Relation to Spread of an Epileptic Attack. *Archives of Neurology and Psychiatry*, *Vol. 4*(44), 740–759.

Williams, B. A. O. (1978). *Descartes: The Project of Pure Enquiry*. Hassocks: Harvester Press.