

Unity of Consciousness: Advertisement for a Leibnizian View ¹

Farid Masrour
Harvard University

It is common to hold that our conscious experiences at a single moment are often unified. But when consciousness is unified, what are the fundamental facts in virtue of which it is unified? On some accounts of the unity of consciousness, the most fundamental fact that grounds unity is a form of singularity or oneness. I call these accounts Newtonian accounts of unity because of their similarity to Newtonian views of space according to which the most fundamental fact that grounds relations of co-spatiality between various points (or regions) of a space is the fact that these points (or regions) are parts of the *same single space*.

It is not, however, clear that the unity of consciousness has to be treated in a Newtonian manner. We can imagine an approach to unity that accounts for it in the same manner that one might think of the unity of a chain. Two links together make a chain when they are connected in the right way. Intuitively, the connection between the links is the fact that grounds the oneness of the chain.

In this paper, I will sketch and defend an analogous account of unity of consciousness. Very roughly, according to this view experiences are unified when they are connected in the right way. In this respect, the view is analogous to Leibnizian views of space according to which the oneness of space emerges from certain conditions over spatial relations. I call this view the connectivity view.

¹ This paper has benefited from discussion with Selim Berker, Ned Block, Matthew Boyle, David Chalmers, Güven Güzeldere, Michael Murez, Efrain Lazos, Christopher Peacocke, Axel Seemann, and Susanna Siegel. More than all, I am indebted to Tim Bayne whose excellent work on unity has influenced my views about the topic.

Section I elaborates on the task at hand, the terminology that the paper relies on, and surveys some of the existing accounts. Section II outlines the main theses of the connectivity view. Section III discusses the comparative dialectical advantages of the view. And the last section addresses three potential objections. I end with a few words about the significance of this issue for cognitive science.

I. Unity and the Grounding Question

As I'm writing these lines, I see them appear on the screen, which partly blocks my view of the window behind it. I feel the pressure of the backrest and a slight pain in my neck. I smell the stew cooking downstairs in the kitchen and hear the cars passing in the street. I can hear the singing of the birds outside. The spring has finally arrived and I feel excited about it. There is something that it is like seeing the screen, hearing the birds, and smelling the stew. But there is also something that it is like for me to have all of these experiences together. Let us refer to this togetherness of my experiences as phenomenal unity of consciousness. Phenomenal unity of consciousness—or, for short, phenomenal unity—is the main topic of this paper.

The above snippet talks about events such as seeing, hearing, and smelling. These events are assumed to be phenomenally conscious events, marked by the fact that there is something that it is like to undergo them. The snippet claims that these conscious events enjoy a form of “phenomenal togetherness.” The “phenomenal” in “phenomenal togetherness” is meant to indicate that there is something additional that this togetherness contributes to what it is like to undergo these experiences, something the omission of which would render our description of the phenomenal facts incomplete.

Here, I have followed the common practice of introducing phenomenal unity with a phenomenological snippet followed by a few clarificatory remarks.² The snippet aims at pointing in the direction of our target phenomenon. Whether this attempt to identify a phenomenon succeeds is itself a matter of controversy and some theorists are skeptical about it. I actually think that there is something correct about this skeptical attitude. But I would like to ask the reader to put any skepticism about the existence of our target phenomenon aside for the moment and regard our characterization as sufficient to serve as a starting point. For reasons that will become clearer in the next two sections, the task of identifying phenomenal unity is closely intertwined with offering a substantive metaphysical account of it. Near the end of section two, we will revisit our initial characterization.

The topic of phenomenal unity has received significant attention during the past two decades.³ Much of the recent discussion has centered around a number of tasks and issues. One task is to provide an account of the most fundamental personal-level facts, if

² See Tye 2003: p. xii, and Bayne 2010: pp. 4-11.

³ See Bayne 2010, Bayne and Chalmers 2003, Dainton 2006, Hurley 2002, Lockwood 1989, Peacocke 1994, Peacocke 2014, and Tye 2003.

any, in virtue of which phenomenal unity obtains (more on this soon.) A second task is to provide an account of the psychological underpinnings of phenomenal unity. A related issue is whether phenomenal unity obtains by necessity or whether it breaks apart in some normal or pathological conditions. There has also been some interest in the relationship between phenomenal unity and personal identity.⁴ Finally, some theorists think that the notion of phenomenal unity is not clearly articulated and that the existence of such a relation has not been properly demonstrated.⁵ Another issue is whether one can provide a satisfactory response to this skeptical worry and if so what the response might look like.

The central question of this paper is about the first task in the above list. I want to answer the following question:

Grounding Question

When several experiences are phenomenally unified, what are the most fundamental personal-level facts, if any, in virtue of which the experiences are unified?

As I am using the term, personal-level facts are those that are present to consciousness. I do not offer a definition of what it is for something to be present to consciousness. But I shall assume that facts about the phenomenal character of experiences and the content of conscious states satisfy this demand. Since personal-level facts must be present to consciousness, the answers to the grounding question can only appeal to facts that are present to consciousness. For example, the answer that a set of experiences is phenomenally unified in virtue of the fact that its members are parts of one encompassing experience satisfies the requirement. Sub-personal underpinnings of conscious states that are not available to the subjects of the states, on the other hand, would not qualify. For example, an answer to the effect that a set of experiences is unified in virtue of the fact that their neural correlates are part of the same corticothalamic loop does not satisfy this demand.⁶ I take the claim that a set of facts obtains in virtue of another set of facts as equivalent to the claim that the former set is grounded in or depends on the other set.⁷

⁴ See Peacocke 2014: ch. 3, Bayne 2010: ch. 12, and Tye 2003: ch. 6.

⁵ See Hill (forthcoming) and Hurley 1998. Hurley does not describe her own position as a form of skepticism. But she defends the position that subjective accounts of unity of consciousness are all bound to fail. Her position implies that unity should not be described in phenomenal terms. She would thus be skeptical about our target phenomenon.

⁶ This, of course, does not mean that there are no difficult cases. Difficulties can emerge in two ways. First, there can be disagreement about how presence to consciousness should be characterized. For example, is introspectability a necessary condition for presence to consciousness? Even if we agree on the requirements for presence to consciousness, there can be disagreement about whether specific cases satisfy this requirement. For example, those who agree on an introspectability criterion might disagree about whether they can find a self or a subject under introspection.

⁷ The cogency and legitimacy of the use of the notion of grounding and its surrounding notions in formulating metaphysical positions has received ample defense in recent literature. For some of the recent contributions to the issue see Fine 2010, 2012a, 2012b, Rosen 2010, Shaffer 2009 and the papers in Correia & Schnieder 2012.

The grounding question should be distinguished from a related question that we can call the structural question: What are the necessary and sufficient personal-level conditions for phenomenal unity? Answers to the grounding and structural questions can diverge. For example, one might give a primitivist answer to the grounding question, holding that phenomenal unity is a fundamental personal-level fact that is not grounded in any other facts at the same level, while answering the structural question by holding that a set of experiences are unified if and only if they are parts of one single encompassing experience.⁸

It is worth noting that the grounding question is posed in terms of unity relations among experiences. My usage of the term “experience” is regulated by Nagel’s criterion.⁹ Accordingly, we are entitled to attribute an experience to a subject, S, whenever we are entitled to hold that there is something that it like to be S. This is not to take a substantive metaphysical position toward experiences, e.g., that they are mental entities such as sense data. I take talk about experiences to be neutral with respect to the debates between representationalist, adverbialist, naïve realist, and sense datum accounts of experience.¹⁰

Formulating the grounding question in terms of experiences might sound problematic to those theorists who hold that we undergo only one experience at each moment. On this one-experience view, one does not have an experience of seeing the screen, an experience of hearing the birds, and yet another experience of smelling the stew. Rather, one has a single experience whose content is only incompletely described in saying that one hears the birds.¹¹ It might seem that if the one-experience view is true, then there is no question to be asked about facts underlying the unity of experiences, because there is no multiplicity in experience. But I think that the challenge that the one-experience view poses for the topic at hand is less substantive than it might seem. In order to pose the grounding question we need elements that have an intimate connection with phenomenal consciousness, are multiple, and stand in unity relations. If the one-experience view is correct, then experience is not the right item to play this role because it fails the multiplicity condition. However, the grounding question can still be posed in terms of contents or their components. One might wonder: when several phenomenally conscious contents together form a unified experience, what is it in virtue of which they do so? I

⁸ The structural claim that a set of experiences are unified if and only if they are parts of one single encompassing experience is compatible with at least three ways to think about the relationship between the two sides of the biconditional. On one view, unity between the experiences is primitive and grounds the existence of the single encompassing experience and the fact that these experiences are parts of the single encompassing experience. On another view, it is the other way around and mereological facts ground unity. On a third possible view, the two sides ground each other and neither is more fundamental than the other.

⁹ Nagel 1974.

¹⁰ There are other difficult issues about the individuation of experiences. For example, is it possible for a subject to have two different tokens of the same type of experience at the same moment in time? The response to questions like this partly depends on how we individuate experiences and there are several options here. However, I do not think that my argument in the paper will be affected by one’s position on this issue.

¹¹ Tye 2003 argues on this basis that there is something problematic about posing the unity question in terms of experiences. See Bayne 2010: pp. 21-28 for a reply to Tye.

therefore think that those who submit to the one-experience view still face the grounding question, though they have to reformulate it in terms of contents. So, although I do not think that the one-experience view is correct, I don't think we need to show that the one experience view is incorrect in order to pose the question about phenomenal unity. In what follows, I will stick to the formulation in terms of experiences.

Another point to note is that the grounding question does not mention subjects of experience. This raises an important question: does the grounding question disappear if the metaphysical structure of an experience essentially involves a subject?¹² No. The idea that each experience involves a subject is compatible with the multiplicity of such subjects. It does not follow from the essential subject-involving nature of experiences that the subject of the act of seeing the screen and the subject of the act of hearing the birds are the same subjects.¹³

A last point to note is that the grounding question does not explicitly mention the temporal relationship between the unified experiences. This diverges from the common practice of separating issues of synchronic unity from diachronic unity. I take it that there is an open question whether a unified account of synchronic and diachronic unities is possible. I do not therefore think that the project of accounting for phenomenal unity should be restricted to synchronic unities. Indeed, we will see later that the connectivity view provides a unified account of synchronic and diachronic unities.

Most recent answers to the Grounding Question are Newtonian. Here are three examples:

- A set of synchronic experiences is phenomenally unified in virtue of the fact that its members are parts of the same single experience. (**mereological view**)
- A set of synchronic experiences is phenomenally unified in virtue of the fact that the contents of its members are parts of the same single content. (**one-content view**)
- A set of synchronic experiences is phenomenally unified in virtue of the fact that its members are experiences of the same single subject. (**one-subject view**)¹⁴

The mereological view grounds phenomenal unity in facts about the oneness of a total experience. Bayne (2010) offers a version of the mereological view.¹⁵ The content view grounds phenomenal unity in the oneness of a content. Tye (2003) advocates a view like

¹² See Peacocke 2014 for a defense of the subject-involving nature of experience.

¹³ Hurley 1998: pp. 99-102 makes a similar claim.

¹⁴ Note that the above views can be co-extensive with each other. For example, it might be the case that the members of a set of experiences are parts of the same single experience iff their contents are parts of the same single content. Nevertheless the views disagree on which one of the sides of the biconditional is metaphysically more fundamental than the other.

¹⁵ See also Bayne and Chalmers 2003.

this.¹⁶ The one-subject view grounds unity in the oneness of a subject. Peacocke (2014) seems to offer a view like this.¹⁷

These Newtonian views are not primitivist views of unity, because they ground it in oneness or singularity. But they are primitivists about one-ness. And it is natural to wonder whether one can provide an account of phenomenal unity that is more substantive than grounding it on a primitive one-ness. My aim in the next sections is to show that this can be done by sketching and motivating a novel Leibnizian view that builds unity from the ground up without an appeal to a primitive oneness. As we shall see, the view also sheds new light on some of the other issues surrounding phenomenal unity in the above list. Thus, those who are primarily interested in the other problems of unity will find some interest in the view that I offer.

2. The Connectivity View

The previous section distinguished between three different accounts of unity namely, the mereological view, the one-content view, and the one-subject view. As we saw, despite their differences, a common Newtonian thread runs through these three accounts. All of these accounts ground phenomenal unity in a global oneness or singularity. My main purpose in this section is to sketch an alternative view of unity that turns the metaphysical order of grounding upside down. On this view, the global unity of experience is grounded in local connections among experiences. The view, therefore, has a Leibnizian structure. As we shall see, this is not the only difference between the connectivity view and its Newtonian rivals. There is a second and equally important difference. But we need to do some stage setting before getting to that.

We can start by highlighting a type of experience that is a common presence in our everyday experiential life. Right now, I see my hand and I see the keyboard. But I also see my hand as being on the keyboard. The experience of my hand as on the keyboard is an experience of a specific spatial relation between them. At the moment, I also hear the

¹⁶ As noted earlier, Tye holds the one-experience view. So he would not formulate his view in the way that I have formulated the one-content view. But as I noted earlier, the grounding question does not disappear by adopting the one-experience view. Tye thus associates the one-ness of an experience with the oneness of content and closure under conjunction. He does not explicitly distinguish between grounding and structural questions. But the manner in which he presents his views suggest that on his view the fundamental fact here is the oneness of content and the closure under conjunction is only a necessary condition for unity.

¹⁷ I am not entirely confident about interpreting Peacocke as giving an answer to the grounding question. The main reason for this is that it is not completely clear to me that Peacocke would regard the one-ness of a subject as something that is present to consciousness. For, he seems to accept the Humean view that we cannot attend to the self and it seems plausible that we would have been able to attend to the self if the oneness of a subject was present to consciousness. If this observation is correct, then Peacocke might be holding the view that the one-ness of the subject is a sub-personal ground for phenomenal unity and phenomenal unity does not have a ground at the personal level. If so, Peacocke's view is a primitivist view about phenomenal unity.

birds singing and feel elated. But that's not all: I experience the singing as the cause of my elation. My experience of the singing as the cause of my elation is another example of an experience of a specific relation.

In my view, the repertoire of experiences of specific relations that we can have is very rich. For example, we can experience spatial, temporal, causal, dynamical, objectual, intentional, and even rational relations. We can experience objects as occupying the same space or as being in more specific spatial relation with each other. We can experience temporal simultaneity and succession relations. We can experience one event as the cause of another. We can experience two events as unfolding in a lawlike dynamic relation with each other. We can experience properties as properties of the same object (this is what I call an objectual relation).¹⁸ We can experience our thoughts and emotions as directed toward objects. And we can experience thoughts or beliefs as based on or justified by the objects and events that we experience.¹⁹ Experiences of specific relations play a central role in the connectivity view.

The intuition behind the connectivity view is that we can account for phenomenal unity in terms of experiences of specific relations. This is partly motivated by the observation that experiences of specific relations seem to suffice for unity. On the connectivity view, the experience of my hand as being on the keyboard is all that it takes to unify my experience of my hand with my experience of the keyboard. The core idea of an account of unity does not need to go beyond this mundane observation. What follows elaborates on this core intuition.

We can use experiences of specific relations to define a relation that we can call binding:

Binding

Two experiences are bound together iff they are connected by an experience of a specific relation.²⁰

My experience of my hand and my experience of the keyboard are bound together by the experience of the spatial relation between my hand and the keyboard. My experience of the singing and my experience of my elation are bound together by the experience of the causal relation between them. In each case, these experiences are connected by an experience of a specific relation. We can call the relation between an experience of a specific relation and the experiences that it binds attachment. Attachment on the connectivity view is a primitive relation.

Binding is a relation among experiences, but bindings are not experienced as relations among experiences. They are experiences of relations among the objects and properties that experiences present. So the idea that there are binding relations among experiences

¹⁸ I borrow this term from Bayne and Chalmers 2003.

¹⁹ All of these claims are substantive and can be opposed. But this paper is not the place to defend them.

²⁰ It is worth pointing out that this binding relation is not the same as the feature binding relation that is discussed in psychology.

does not conflict with the diaphanousness or transparency of experience.

It seems plausible that two experiences can be unified without being bound. Take the earlier example in which you see your hand on the keyboard and hear the singing of the birds as the cause of your elation. Let us add that you hear the singing as coming from a different region in the same space in which you experience the keyboard to be located. Then you are experiencing a spatial relation between the keyboard and the singing. So your visual experience of the keyboard and the auditory experience of the singing are bound together. More importantly, it seems plausible that your experience of the keyboard and your experience of the elation are unified because they are both bound to the experience of the singing. So your experience of the keyboard and your elation experience are unified but not bound.

One upshot of this observation is that the intuitive notion of unity is weaker than the binding relation. Intuitively, two experiences that are bound are unified with each other. But not all experiences that are unified need to be bound with each other. Therefore, we should not identify unity with binding. Nevertheless, there can be an intimate connection between the two notions. The above example also illustrates this connection. The seeing of the keyboard and the experience of the elation are connected through the mediation of the experience of the singing that is bound to both. There is, as it were, a path between these two experiences that is instantiated in virtue of a chain of bindings. More precisely, there is a unity path that connects these two experiences, where a unity path is defined as follows:

Unity Paths

There is a unity path between two experiences E_m and E_n iff E_m is bound with E_n or there is an E_r such that E_m is bound with E_r and there is a unity path from E_r to E_n .

A path consists in a chain of binder experiences (experiences of specific relations) and the experiences that are bound by the path. An experience can be a member of a unity path in either of two ways:

Path Membership

An experience is a member of a path iff it is one of the binders in the path or one of the experiences that are bound by the binders in the path.

The notions of a unity path and path membership can be used to define a property of a set of experience that I call minimal connectivity:

Minimal Connectivity

A set of experiences, S , is minimally connected iff there is a unity path, P , such that all the members of S are members of P .

We are now in the position to characterize the central claim of the connectivity view. I call this claim the connectivity thesis:

Connectivity Thesis

A set of experiences is unified in virtue of the fact that it is minimally connected.

It is worth emphasizing the asymmetric relation between unity and connectivity. According to the connectivity thesis, unity is grounded in minimal connectivity. The connectivity thesis is thus an answer to the grounding question. The structural analog of the connectivity thesis would have been: a set of experiences is unified iff it is minimally connected. This biconditional thesis is silent about issues of metaphysical priority.

On the connectivity view unity relations are grounded in the existence of unity paths and facts about membership in the unity path. One way to think about the grounding relation between unity and unity paths is to think of unity paths as determinate versions of the determinable unity. Intuitively, determinables are instantiated in virtue of the instantiation of their determinate versions.²¹ This idea is not essential to the connectivity view, but provides an additional framework for understanding it.

The connectivity view is Leibnizian in the sense that it accounts for unity in terms of local relations. This distinguishes this account from the Newtonian views that we discussed in the previous section. But there is another equally important way in which the account diverges from the existing accounts. To bring out this difference, let us consider the following paragraph from Bayne & Chalmers (2003):

... when I look at the book while feeling a pain, there is something it is like to see the book (yielding a phenomenal state A), and there is something it is like to feel the pain (yielding a phenomenal state B). But there is more than this: there is something it is like to see the book while feeling the pain. Here there is a sort of conjoint phenomenology, that carries with it the phenomenology of seeing the book, and the phenomenology of feeling the pain. ..., we can think of the conjoint state here as involving at least the conjunction A&B of the original phenomenal states A and B.

On the Bayne-Chalmers view, the fact that the feeling of pain and the seeing the book are unified with each other makes a difference to the overall phenomenology of the subject. But this difference is a matter of seeing the book and feeling the pain in a conjoint manner. On this view having two experiences in a conjoint manner does not make a substantial contribution to phenomenology or content. Unity is a purely structural or logical matter and the phenomenology associated with it is, as it were, bare phenomenology.

This observation generalizes to the other Newtonian views that we considered in the previous section. On all of these views unity is a purely structural feature. We can call these views bare unity views, where bare unities are connections between experiences that can happen independent of any experience of a specific relation.

²¹ See Rosen 2010: pp.126-8 for a defense of this claim.

On the connectivity view, in contrast, phenomenal unity between experiences makes a substantial contribution to overall phenomenology and content. On this view, if I tell you that my experience of the singing of the birds and my feeling elated are bound, there is something left out in my description of my phenomenology. You can still ask: “How are they bound?” And I can give you an informative answer: “They are bound in that I experience the singing as the cause of my elation.” A similar idea applies when experiences are connected but not bound. In such cases too, there is a unity path that connects the experiences and the unity path makes a substantial contribution. The second way in which the connectivity view diverges from the common Newtonian has to do with this contrast. We can say that phenomenal unity on the Newtonian views is a bare relation, but on the connectivity view it is a substantial relation. We will see in the next section that this puts the connectivity view in a dialectically advantageous position in comparison to bare unity views.

The substantiality of unity under the connectivity view is independent from the fact that the view has a Leibnizian structure. One can easily conjure up a Leibnizian view that is structurally similar to the connectivity view but its binding relations are bare unity relations.

This ends my characterization of the core concepts and theses of the connectivity view. On this view, the unity of a set of experiences is a matter of the connectivity of the set. Before ending the section, I want to return to how I introduced phenomenal unity at the beginning of the paper and remove an ambiguity in the introduction. There, I introduced unity as a form of phenomenal togetherness. We should note that the referent of the phrase “togetherness of the experiences” depends on whether we adopt a Newtonian view or the connectivity view. On the Newtonian view, the togetherness refers to a bare global phenomenal oneness. On the connectivity view, in contrast, the togetherness refers to the connectivity of my experiences. This connectivity is implicitly contained in the passage. The assertion that the screen partly blocks my view of the window is meant to imply that I experience a spatial relation between the screen and the window. It is also implicit that the stew that I smell, the birds and the cars that I hear, and the backrest whose pressure I feel on my back are all experienced to be located in the same space in which I feel my body to be located. The togetherness refers to the fact that all of my experiences are connected in this way. So there was an ambiguity in our initial characterization of phenomenal unity that we were not in the position to remove then. Now that we have a clear picture of the contrast between the connectivity view and the Newtonian accounts, we are in a position to remove this ambiguity.

3. The Dialectical Advantages of the Connectivity View

We saw in the previous section that the connectivity view diverges from the Newtonian approaches to unity in two ways. First, the view has a Leibnizian character in that it accounts for phenomenal unity in terms of local relations. Second, on the connectivity view unity relations are substantial and make a positive contribution to the phenomenology and content of the experiences that they unify. These two differences put the connectivity view in a dialectically more advantageous position in comparison to the

Newtonian views. My aim in this section is to defend this claim by arguing that the connectivity view is in a better position in three respects.

Consider the following possible account of phenomenal unity:

The one-stream view

A set of experiences is phenomenally unified in virtue of the fact that its members belong to the same stream of consciousness.

There is something about this view that leaves us cold. The notion of phenomenal unity seems to be too close to the notion of a single stream of consciousness in terms of which it is elucidated.²² Many, therefore, might claim that they grasp the idea that a set of experiences is unified and the idea that the set forms a single stream of consciousness in a similar way. In short, the explanandum and the explanans in the one-stream view are too close to each other. This is not to say that the account does not tell us anything substantive. The account gives us a Newtonian picture of unity which is substantively different from a Leibnizian picture. But there is an intuitive sense in which the account does not expand our knowledge of the matter. To use a Kantian term in a slightly different manner, the account is not ampliative.

Most Newtonian accounts suffer from a similar dialectical shortcoming, though to different degrees. Consider, for example, the mereological account according to which experiences are unified in virtue of the fact that they are parts of the same single experience. The account does not seem to lack substance. It is, after all, a Newtonian account according to which a oneness grounds unity. The claim that what makes consciousness unified is this Newtonian structure seems to be substantive. Also, the claim that the unifying oneness is an experience is, on the face of it, a substantive claim. A mountain range is a unified whole whose parts are mountains. But a mountain range itself is not a mountain. In claiming that a unified set of experiences together form an experience the account might be making an additional substantive claim.

But under closer inspection, the mereological view suffers from a shortcoming that is similar to the one-stream view. Consider a view about the individuation of experiences that we might call object-based individuation. On this view, we should cut up the space of experiences in the same way that we cut up the space of their objects. Under this object-based notion of experience, I have an experience of the keyboard and an experience of elation. But I do not have an experience of the keyboard and elation because my elation and the keyboard do not together form an object of experience. It is not obvious that the object-based notion of experience is correct. Nevertheless, we seem to have an intuitive grasp of the object-based notion of experience. But obviously, the notion of a single encompassing experience that is operative in the mereological view is a different notion. For, if we adopt the object-based view, many of the experiences that we often have at the same moment cannot be regarded as forming one single experience together. This gives rise to a question: what is the notion of a single experience that is operative in the

²² Van Gulick 2013 notices a similar point about Bayne's mereological account but puts it to a different use.

mereological view and is our grasp of this notion sufficiently removed from our grasp of a unified set of experiences? In my view, the notion of a single experience that is operative in this account is too close to the notion of unity and our grasp of this notion is not independent from our grasp of the notion of unity. If this observation is correct, then the mereological view suffers from a similar shortcoming as the one-stream view. Its explanandum and its explanans are not sufficiently distant from each other. Therefore the account is not sufficiently ampliative.²³

The situation for the other Newtonian approaches seems analogous. On the one-content view, experiences are unified in virtue of the fact that their contents are parts of the same total phenomenal content. The view does not ground unity in the one-ness of experience. Rather, it grounds it in the oneness of phenomenal content. But it is not clear whether we can have a clear grasp of the oneness of phenomenal content independent of our grasp of unity. Again, the explananda and the explanans are too close to each other. Arguably, similar issues arise for the one-subject view.²⁴ All Newtonian views thus suffer from the same dialectical shortcoming. We can say that these accounts are not very ampliative.²⁵

The first dialectical advantage of the connectivity view is that it is ampliative. This advantage emerges out of the interplay between its Leibnizian structure and its substantiality—the fact that under the connectivity view unity is not a bare relation. The connectivity view grounds phenomenal unity on relations of binding and those in turn on the attachment between experiences of specific relations and other experiences. The relations of binding and attachment are clearly different from the unity that the account aims at explaining. The account's explananda and explanandum are clearly distinct. Thus the connectivity view is clearly ampliative and substantive. In this respect it is dialectically preferable to the Newtonian views that we have considered.

The second dialectical advantage of the connectivity view has to do with the fact that those who wish to provide an account of phenomenal unity have to defend their claim that there is a target phenomenon to be accounted for. Here, the dominant Newtonian accounts of unity have met some skepticism. For example, in a commentary on Bayne's *Unity of Consciousness* Hill complains that “[Bayne] has not yet specified an appropriate relation of phenomenal unity, or even pointed us in a direction in which an appropriate unity relation can be found.”²⁶ In the book, Bayne claims that introspection supports the existence of phenomenal unity. But Hill responds that he can never find Bayne's

²³ Bayne 2010 offers what he calls a tri-partite conception of experiences according to which experiences are individuated on the basis of their subjects, their phenomenal properties, and their time of occurrence. In my view, this account pushes the question of individuating experiences to the question of individuating phenomenal properties and that is an issue about which we may not have clear intuitions.

²⁴ For an argument to the contrary of this see Peacocke 2014: ch. 3.

²⁵ The claim that an account has a dialectical shortcoming because it is non-ampliative should be understood in the context of the availability of competing more ampliative and equally plausible accounts of the same phenomenon. In such a context, we have some reason to abandon the non-ampliative account in favor of the ampliative one. This is not to say that this reason cannot be overridden by other considerations.

²⁶ Hill (forthcoming).

phenomenal unity whenever he introspects, “looking for a phenomenal unity relation.” So the claim that there is phenomenal unity has been challenged and a well-developed account of unity has to respond to this challenge.

In my view, the main reason that Bayne’s view has received a skeptical reaction of this sort is due to a feature that it shares with all other Newtonian accounts. As we noted earlier, in grounding unity in a form of oneness, Newtonian accounts are forced to construe it as a bare relation. The connectivity view, in contrast, grounds phenomenal unity in experiences of specific relations. It is not as easy to be skeptical about substantial experiences of specific relations as it is easy to be about bare phenomenal unities.²⁷ Thus the connectivity view is in a better position to convert some of the skeptics about phenomenal unity into friends of phenomenal unity.

The third, and perhaps the most important, dialectical advantage of the connectivity view is that it provides a unified account of diachronic and synchronic phenomenal unity. On this view, the conditions for diachronic unity between experiences are exactly the same as the conditions for synchronic unity: the existence of unity paths between them. Proponents of Newtonian views, in contrast, often admit that their account of synchronic unity does not generalize to diachronic unities.²⁸ Clearly, a uniform account of unity is theoretically preferable to a non-uniform one in that it accounts for more phenomena in a uniform manner.

This ends my defense of the claim that, in three respects, the connectivity view is in a dialectically more advantageous position than the existing Newtonian accounts. First, the view is ampliative because we grasp its explanandum and explanans differently. Second, the view is in a better position to face the skeptical challenge about phenomenal unity because unity relations on this view are not bare. Third, the view is theoretically in a better position because it provides a unified treatment of synchronic and diachronic unities. I think it is safe to conclude that we have good reasons to embrace the connectivity view unless there are important objections to it. Whether there are such objections is the focus of the next section.

4. Objections and Replies

The previous section developed the connectivity view and argued that it is in a dialectically better position than its rivals. My aim in this section is to block three possible objections to the view.

Hurley has argued against attempts to account for unity in subjective terms by employing a generic argument style that she aptly calls the “just-more-content” argument.²⁹ The argument is based on the observation that if a set of experiences is not unified, adding

²⁷ Hill, for example, can find some forms of unity like spatial unity under introspection. His problem is with finding a unity relation that obtains universally.

²⁸ See Bayne 2010, Peacocke 2014, and Tye 2003.

²⁹ Hurley 1998: pp. 97-102.

more content can unify its members only if the new content is already unified with them. This, Hurley concludes, shows that we cannot account for unity in terms of the contents of experience. I am not concerned with the cogency of this argument. But I would like to consider a possible objection that one who finds Hurley's argument attractive might submit against the connectivity view. Our imaginary interlocutor might reason in the following manner:

Imagine a situation in which a subject has two streams of consciousness. In one stream the subject has three experiences: E1, E2, and E3. Let us assume that E3 is an experience of a specific relation between the contents of E1 and E2 and binds the two experiences together. In the other stream the subject has an experience, E4, which is type-identical with E2. Since E2 and E4 have the same content and E3 is an experience of a specific relation between the contents of E1 and E2, it follows that E3 is also an experience of a specific relation between the contents of E1 and E4. Under the connectivity view then, it follows that E1 and E4 are bound and unified. However, E4 is in a different stream from E1 and E2, which implies that E4 is not unified with E1 because by stipulation, it is in another stream. Thus the connectivity view results in a contradiction and should be rejected.

This is an interesting objection, but it does not survive scrutiny. Consider the thesis that a subject's two experiences, E1 and E2, with contents p and q are bound in virtue of the fact that the same subject has an experience with the content R(p, q), where R(p, q) is the content that an experience type that would bind E1 and E2 must have.³⁰ This thesis has many counterexamples and the above scenario illustrates one of them. But the thesis is different from the connectivity view under which two experiences are bound when an experience with the appropriate content connects them together. Under the connectivity view, having an experience with the appropriate content does not suffice for binding. The experience must be attached to the experiences that it binds. So the connectivity view is compatible with the above scenario that illustrates a situation in which the instantiation of the sufficient content in the subject does not suffice for binding. The connectivity view is not an attempt to ground unity on content.

But our imaginary interlocutor might respond:

But there should be an explanation for the fact that E3 binds E1 and E2, but does not bind E1 and E4. And the best explanation is that E4 is not in the same stream of consciousness as E1 and E3. So experiences of relation can bind other experiences only when the binder and the experiences that it binds happen within a single unified stream. Thus what partially grounds the fact that the appropriate content binds two experiences is that the binding experience and the experiences that are bound are unified with each other. It therefore follows that unity is presupposed by binding and cannot be grounded in it.

³⁰ Here, I am using the brain-body notion of a subject under which there is one subject when there is one brain-body. Under this sense, there is no contradiction in assuming that one subject has two streams of consciousness.

The above move extends the intuition behind the just-more-content argument. It is admitted on both sides that adding more content does not unify two experiences. But our imaginary interlocutor takes this a step further by arguing that the best explanation for this is that the additional content is not unified with what it aims to unify.

However, on the connectivity view attachment is a primitive relation in the sense that there is no other fact at the personal level that grounds it. E3 binds E1 with E2 but not E1 with E4 in virtue of the fact that E3 is attached to E2 but not to E4. But this fact has no further ground at the personal level that metaphysically explains it. In assuming the contrary, the above argument is begging the question against the connectivity view. Let us move to the second objection.

The second possible objection to the connectivity view concerns its breadth:

It seems that the connectivity view cannot account for all cases of phenomenal unity. For example, my emotional experiences are unified with my perceptual experiences but I don't find any connection between them. A similar point applies to the case of experiences associated with thoughts and other cognitive states. It is not clear how these experiences might be connected to my other experiences through unity paths.

Since the objection argues that the connectivity view does not have the sufficient resources to account for the unity between all of our experiences, I call it the insufficiency objection. Does the insufficiency objection succeed?

The first thing to note is that a number of different theoretical options are available about the nature of emotional experiences. Under one view, emotional experiences are collections of experiences of bodily dispositions. We can experience a variety of bodily dispositions. For example, we can sometimes feel disposed to laugh, to cry, to embrace, to jump, to dance, to sit, to run, to shout, to punch, and so on.³¹ On a view like this, there is something that it is like to feel joyful at each moment. But when we make this claim we are not adding anything to the repertoire of the simple experiences that we can have because feeling joyful is a label for some collection of experiences of bodily dispositions such as feeling disposed to jump, dance, and sing. On a different view, feeling joyful consists in undergoing an unanalyzable raw feel of joy. On a third view, feeling joyful essentially combines a raw feel of joy with experiences of bodily dispositions. Obviously, the way in which emotional experiences can connect to perceptual experiences depends on which one of these views one takes. For example, if experiences of emotions essentially involve experiences of bodily dispositions then it is not hard to imagine how they can connect to perceptual experiences.

Even under the view that experiences of emotions are simple raw feels, there are several ways in which these experiences can connect to perceptual experiences. The most important potential source of connection is the experience of temporal relations. One can

³¹ These claims are of course substantive and require proper defense. But this paper is not the place to do so.

experience one's emotion to be simultaneous with what one's perceptual experience presents. One can also experience one's emotions as temporally succeeding or preceding what one's perceptual experience presents.³² Also, emotions can be experienced as standing in causal relations with external objects and events as well as our behavior, thoughts, and memories. For example, one might feel the sight of the snake as the cause of one's fear or one might feel one's sadness as the cause of one's crying. We might feel our happiness as caused by remembering a past joyful event. We can also experience our emotions to be directed at the items that perceptual experiences present. For example, one might experience one's love as directed towards an individual. We can feel our emotions to be epistemically related to certain judgments and beliefs about oneself. For example, a sincere judgment "I'm excited" can be felt to be based on one's excitement.

Thus, the proponent of the connectivity view has several resources to account for the connection between emotional and perceptual experiences. First, she might defend an account of the nature of emotional experience that directly links it to bodily dispositions and through them to other perceptual experiences. Second, she might hold that experiences of temporal, causal, intentional, and epistemic relations bind emotional and perceptual experiences. Finally, she might find connections between emotions and perceptual experiences through their connections with behavior, thoughts, imaginings and memories. Arguably, a similar line of response would work for the experiences associated with cognitive states. So we can safely conclude that the insufficiency objection can be resisted. Let us thus turn to the last potential objection.

On the connectivity view, a set of experiences cannot be unified unless it is minimally connected. The third potential objection against the connectivity view targets this idea. The objection attempts to show that unified sets of experiences that are not minimally connected are possible. Since the objection bases this claim on the conceivability of such situations, I call it the conceivability objection. Here is one way that the objection might go:

I find unified sets of experiences that are not minimally connected conceivable. For example, I can conceive of an emotional experience that has no connection with my other experiences but yet is unified with them because it is experienced from the same phenomenal point of view. In general, one can conceive of experiences that are disconnected but unified because their contents are experienced as given to a single point of view.

It is not entirely clear what a phenomenal point of view is. But whatever a phenomenal point of view might be, the natural position for the proponent of the connectivity view is that the oneness of a phenomenal point of view is grounded in phenomenal unity. So on the connectivity view a situation in which an emotional experience is completely

³² In my view, simultaneously experiencing two things is not an experience of simultaneity between them. Neither is successively experiencing two events the experience of succession. So it is not the case that what synchronic experiences present are always experienced as simultaneous and what successive experiences present are always experienced as successive.

disconnected from all other experiences yet experienced from the same phenomenal point of view is impossible. But can we conceive of such a situation?

I have been emphasizing that the proponent of the connectivity view has several resources to account for the connection between emotional experiences and other experiences. In order to conceive of an emotion that is disconnected from other experiences we have to conceive of a situation in which all potential sources of connection are absent. And it is not clear to me that after making the absence of all of these connections explicit, we would still confidently assert that the alleged emotion is given to the same point of view as our other experiences. Let me elaborate

Typically, when I conceive of an emotion along with other experiences, I conceive of a situation in which I experience the emotion as simultaneous with what my other experiences present: e.g., I felt my fear as I witnessed the biker zigzagging between the cars.³³ I can also conceive of an emotion as preceding or succeeding my perceptual experiences. In order to conceive of a fully disconnected emotion I have to conceive of a situation in which these temporal connections are absent. Typically, when I conceive of an emotion I also conceive of a situation in which I experience some bodily dispositions. For example, when I conceive of feeling excited, I conceive of being disposed to get up and pace back and forth. In order to conceive of a fully disconnected emotion I have to conceive of the absence of all of these dispositions. It is not completely clear to me that I can conceive of a situation like this. More importantly, the further that I go down the path of imaginatively cutting down the connections between an emotion and other experiences, the less confident I feel that I can imagine the emotion to be given to the same point of view as my other experiences.

So the gist of my reply to the conceivability objection is that if fully conceiving of a situation requires conceiving of all the relevant detail in that situation, then it is not at all clear that we can fully conceive of a situation in which an emotional experience is fully disconnected from our other experiences, yet given to the same point of view. Admittedly, more work needs to be done on what the oneness of a phenomenal point of view is and in virtue of what it emerges. But I think that when we understand this aspect of phenomenal experience in a deeper way, the *prima facie* conceivability of the situation that the conceivability objection depicts disappears.

To summarize, I have considered three possible objections to the connectivity view. On the just-more-content objection, the fact that an experience of a specific relation can bind

³³ Earlier, I said that experiences of relation are not experiences of relations between experiences, but experiences of relation between what my experiences represent. One might think that this would make extending the view to emotional experiences or experiences such as moods somewhat difficult. For, it is not completely clear what these experiences represent. However, I do not think that this is a serious worry about the connectivity view. The remark that all experiences of relations are experiences of relations among what experiences represent is presupposing a representationalist view of all experiences. If it turns out that representationalism cannot be extended to some experiences, then we have to say that some experiences of relations are experiences of relations between experiences. Nothing in the connectivity view would change as a result.

two other experiences is itself grounded on the fact that the binding experience is unified with what it binds. I responded that this claim begs the question against the thesis that relations of attachment are primitive relations. On the insufficiency objection, the connectivity view cannot account for the unity between emotional and cognitive experiences with other experiences. In response, I pointed out several ways in which these experiences can connect to other experiences. Lastly, I considered the objection that fully disconnected and yet unified experiences are conceivable and thus possible. I responded by questioning the claim that such situations are conceivable.

Conclusion

My aim in this paper has been to advertise for a novel account of phenomenal unity. To this effect, I outlined a view according to which a set of experiences is unified when there is a unity path that contains all of its members. On this view unity paths are chains of binding relations, which are mediated by experiences of specific relations. I called the view the connectivity view. We saw that the connectivity view is different from the common contemporary approaches to phenomenal unity in two respects. First, unlike the Newtonian views, the connectivity view has a Leibnizian structure. Newtonian views ground phenomenal unity in a primitive form of oneness. The connectivity view, in contrast, builds unity from local unity relations and does not appeal to a primitive oneness. I then argued that the connectivity view is dialectically in a better position than its Newtonian rivals in that (a) the view is more substantive, (b) it is in a better position to face the skeptical challenge about phenomenal unity, and (c) it treats synchronic and diachronic unities in a uniform fashion. Finally, I responded to three potential objections to the view.

I want to end the paper by quickly noting the implications that the choice between the connectivity view and its Newtonian rivals might have for the two other issues surrounding phenomenal unity—namely, the issue of whether consciousness is in some sense necessarily unified and the issue of the psychological underpinnings of unity.

It is not uncommon for the Newtonian theorists about unity to hold that consciousness is necessarily phenomenally unified.³⁴ Part of the reason for this is that under these views, the failure of unity seems in some sense inconceivable. The connectivity view, in contrast, gives us a better handle on what it takes for unity to break. Phenomenal unity breaks if there is discontinuity in the stream of consciousness. So arguably, one implication of the connectivity view is that the thesis that consciousness is necessarily phenomenally unified is false.

The connectivity view also seems to give us a different picture of the cognitive architecture that underlies unity. On a Newtonian view, unity requires a phenomenal oneness that encompasses all experiences. On the assumption that unity structure is mirrored in cognitive architecture, the view seems to require an architecture in which experiences all belong to a single center. Newtonian views thus suggest, to borrow a label from Bayne, an imperial cognitive architecture. The Leibnizian view, in contrast, is

³⁴ See Bayne 2010 and Bayne & Chalmers 2003.

compatible with a cognitive architecture in which there is no single center to which all experiences belong. The view is more at home with a feudal architecture in which several local centers are minimally connected with each other.³⁵ Thus, if we assume that conscious experiences are often unified and that unity structure is mirrored in cognitive architecture, adopting the connectivity view would have important consequence for the way that we think of the cognitive architecture that underlies our experience.

Works Cited

- Bayne, T., & Chalmers, D. (2003). What is the Unity of Consciousness? In Cleeremans, A. (Ed.), *The Unity of Consciousness: Binding, Integration, and Dissociation*, 23-58.
- Bayne, T. (2010). *The Unity of Consciousness*. Oxford University Press, USA.
- Correia, F., & Schnieder, B. (Eds.). (2012). *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge University Press.
- Dainton, B. (2006). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. Taylor & Francis.
- Fine, K. (2010). Some Puzzles of Ground. *Notre Dame Journal of Formal Logic*, 51(1), 97-118.
- Fine, K. (2012a). Guide to Ground. In Correia, F., & Schnieder, B. (Eds.), *Metaphysical Grounding: Understanding the Structure of Reality*, 37-80. Cambridge University Press.
- Fine, K. (2012b). The Pure Logic of Ground. *The Review of Symbolic Logic*, 5(1), 1-25.
- Hill, C. (forthcoming). Tim Bayne on the Unity of Consciousness. *Analysis Review*.
- Hurley, S. L. (1998). *Consciousness in Action*. Harvard University Press.
- Lockwood, M. (1989). *Mind, Brain and the Quantum*. Oxford: Basil Blackwell.
- Nagel, T. (1974). What Is It Like to Be a Bat? *The Philosophical Review*, 83(4), 435-50.
- Peacocke, C. (Ed.). (1994). *Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind*. British Academy.
- Peacocke, C. (2014). *The Mirror of the World: Subjects, Consciousness, and Self-Consciousness*. Oxford University Press.
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In Hale, B. & Hoffman, A. (Eds.), *Modality: Metaphysics, Logic, and Epistemology*, 109-36. Oxford University Press.

³⁵ I also borrow the term “feudal” from Bayne 2010.

Shaffer J. 2009. "On What Grounds What," In Chalmers, D., Manley, D. & Wasserman, R. (Eds.), *Metametaphysics*, 347-83: Oxford University Press.

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

Van Gulick, R. (2013). Phenomenal Unity, Representation and the Self. *Philosophy and Phenomenological Research*, 86(1), 209-14.