13 Experiences and Their Parts

Geoffrey Lee

1 Introduction

Intuitively, your overall conscious experience at a given time has a complex structure. For example, you might simultaneously have experiences in more than one sensory modality, and your experience within a single modality might involve awareness of a complex of different features of a stimulus, of multiple stimuli, or of relations between stimuli. You might also simultaneously have both perceptual and nonperceptual experiences such as thoughts or conscious emotions; these states have their own structure—for example, a thought might involve using an array of different concepts—and their existence contributes to the overall structure of your conscious state.

Perhaps the simplest kind of structure that experiences may have is a quasi-mereological structure that results from some experiences being components of, or "subsumed" by, more complex experiences. For example, your experience of the color and shape of an object might subsume your experience of its shape. There are arguably other kinds of structure that experiences have—for example, perhaps they have a representational structure that is richer than a hierarchy of quasi-mereological relations such as a sentential structure with elements playing the role that names and predicates play in English, or an imagistic structure (compare Fodor's [1975] view that beliefs and other propositional attitudes have representational structure). A general understanding of experiential structure should be a central goal of a theory of consciousness—for example, it is necessary for understanding how experiences relate in a *systematic* way to neural states. The relatively modest task of giving an account of the part/whole structure of experience is my concern here: we can think of this as a first step toward a more general theory of experiential structure. Moreover, my aim is doubly

modest in that I will mostly taxonomize some different views rather than make a positive proposal.

My discussion will focus on a number of interrelated core questions about experiences and their parts:

The Identification Question: What is the class of "experiences?" What are the *complex* experiences whose part/whole structure we are interested in?
The Parthood Question: When we talk about the "parts" of an experience, what is the relevant parthood relation?
Experiential Decomposition Question: Given an experience, what are the parts of the experience?
Experiential Composition Question: Given two or more experiences, under what conditions are they components of a more complex experience?
The Priority Question: Is a complex experience constructed from its experiential parts, or are they derivative from the whole?
The Unity Question: What is the relationship between the "unity" of consciousness and the part-whole structure of experiences?

My main concern here will be the priority question, although all the questions will be relevant. A number of recent authors draw a distinction between holistic and atomistic views of consciousness (Searle, 2005; Bayne & Chalmers, 2003; Bayne, 2010), but I think it is fair to say that we lack a detailed account of the different ways this distinction can be drawn and how they are related. For example, Searle (2005) contrasts what he calls the "building block" model of consciousness with his preferred model on which consciousness involves a single, unified "field," whose parts are not separately existing events, but rather modifications of the whole field. Searle suggests that the field model has important methodological consequences, in particular that attempts to explain consciousness in a piecemeal way by finding local neural correlates of different conscious contents is wrongheaded because the conscious field is not constructed from such independent parts. The building-block/holistic-field contrast is intuitive, but ideally we would like a rigorous account of what the different distinctions are in this ballpark. I hope to make some progress toward giving such an account here. My main point will be that the building-block metaphor adopted by Searle is misleading because it is suggestive of a number of different theses about the parts of an experience, all of which an anti-holist—understood as someone who thinks that typical experiences *do* have parts that are independent of the whole—could reject. There is a risk of saddling the anti-holist with the view that a total experience at a time can be broken down into a set of completely independent atomic experiences, which are

somehow glued together to produce the complex whole. More precisely, the building-block metaphor is suggestive of four theses: *supplementation, constructivism, disjointness,* and *unity externalism*, which I will discuss in detail, arguing they can all be rejected without accepting the holistic field view.

The bricks and other parts of a building satisfy *supplementation* in the sense that each proper part is accompanied by another disjoint proper part. This further implies[1] that the building can be fully broken down into its proper parts in the sense that any given part of the building is accompanied by a set of mutually disjoint parts that sum to the whole. I will discuss how to formulate analogous claims about the parts of an experience, explain why they are probably not correct, and why rejecting holism doesn't require accepting them. Thus an anti-holist need not think that a total experience can be fully broken down into independent experiential parts.

A building is *constructed* out of a set of disjoint proper parts, in the sense that if the parts exist and stand in certain relations to each other, then we have a building. We might similarly think of a complex experience as "constructed" out of simpler experiences, in virtue of a unifying relation obtaining between them. In this way, we might think that in order to understand experiential composition, we need to figure out what the unifying relation is. As we will see, however, this constructivist approach to complex experiences might be wrongheaded. For example, it might be that experiential wholes are prior to their parts, as holists contend, and we should be starting with the wholes and asking how they *de*compose into parts. I will argue that anti-holists can *also* reject the constructivist approach, treating unity in a "top-down" way. On this view, there may be a sense in which there is no separately specifiable unity relation between experiences.

Finally, bricks are *disjoint* objects, and in order for them to be "unified" into a house, they must stand in certain *external* relations such as being cemented together. I think the anti-holist can reject the claim that the parts of an experience are disjoint—they may overlap in a sense I will explain. Furthermore, it is consistent with rejecting holism that the unity relation between them is an *internal* relation, not an external relation, and so the metaphor of unity as a kind of phenomenal cement might be misleading.

The plan: In section 2 I discuss some preliminary issues regarding the identification question and the parthood question. In section 3 I explain how I think we should understand the distinction between holism and atomism. In section 4 I discuss supplementation, in section 5 I discuss constructivism, and in section 6 I discuss disjointness and unity externalism.

2 Preliminaries: Identification and Parthood

What are the "experiences" whose part/whole structure we are interested in? I will assume that experiences are events, and they involve the instantiation by subjects of certain special properties, call them "experiential properties," such that what it's like to be a subject at a given time is constituted by the experiential properties they enjoy. Thus, experiences are the components of "phenomenal consciousness," as many philosophers call it.

Let's assume that phenomenal consciousness is a reasonably well-defined phenomenon, acknowledging that some may dissent (see, e.g., Byrne, 2009). Simply making this assumption, however, does not answer some basic—and surprisingly problematic—questions about the *extension* of the term "experience" that will be salient here. What are the experiences you are having right now? Are you enjoying an overarching über-experience, a "field of awareness" that encompasses all the more local experiences you're having? Or is this "field" at best a mere sum of more local experiences? And what experiences does an experience have as parts? Does your experience of an object have separate experiences of its features as parts, or are these "experiences" really just aspects of your complex experience of the whole object? Does a multimodal experience decompose neatly into modality-specific parts? Beyond the present moment: do experiences extend through time, perhaps having experiences as proper temporal parts? Or perhaps they are always instantaneous, or at least lacking temporal parts? In short, it is unclear when exactly experiences *compose* other experiences, and it is unclear how exactly experiences *decompose* into other experiences.

On the one hand, it is natural and sensible to suspect that the answers to at least some of these questions will turn on stipulative decisions about how to use the term "experience" and so are not really substantive. On the other hand, experience presumably does have some kind of objective structure that we can discover; if there are verbal issues, they presumably have to do with which parts of this objective structure to call "experiences". Furthermore, it remains true that we need to start with at least some clear instances of conscious experience in order to study it. I'll assume that we can start with *paradigmatic* cases of experiences, such as experiences of features like colors and shapes, object experiences, and experiences of relations between objects such as spatial relations. We can then ask under what circumstances these compose more complex experiences, and how they decompose into simpler experiences. We should note right from the beginning that such "paradigmatic experiences" may not be *self-standing*: that is, their nature and existence might be metaphysically dependent on other experiences

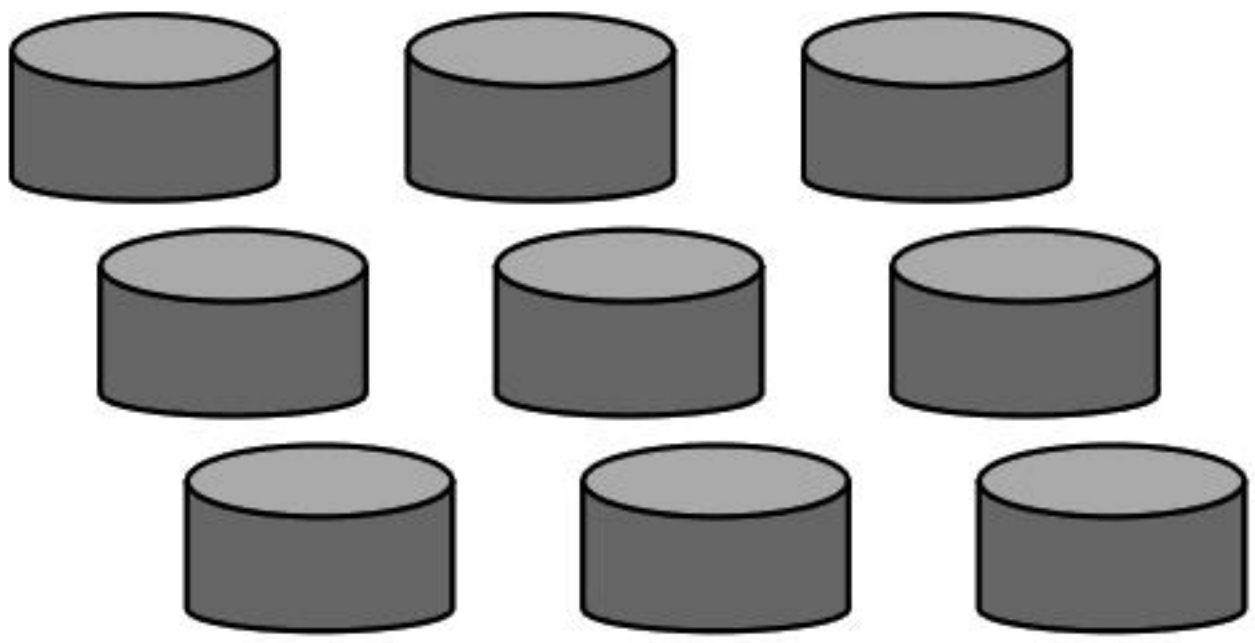

Figure 13.1
Gray cylinders.

(e.g., experiences of color and texture may be interdependent in this way). I think we cannot say in advance what the self-standing experiences are; for example, if holism is true, they may include *only* whole multimodal "fields" of awareness and therefore be quite distinct from the paradigmatic experiences, which may be mere *aspects* of self-standing experiences.[2]

As another preliminary, I want to put on the table a distinction between two approaches to the parthood relation on experiences, the type-entailment approach and the realization approach. So that we have a concrete example to consider, contemplate your experience of figure 13.1. It includes as parts experiences of the individual cylinder shapes, experiences of spatial relations between the cylinders, and experiences of distinct features of each cylinder such as their color and shape. In what sense are these parts of the whole experience?

Talk of "parts" of experiences is puzzling. Is this the same part/whole relation that links bricks with houses, or is it merely some other relation that is sufficiently parthood-like that speaking in this way is excusable? I'm doubtful whether there is a single part/whole relation that applies in every case where we talk in mereological terms, but I will not take a stand on this issue here (for discussion see Bennett, 2011; Fine, 2011; and Uzquiano, 2006). We'll consider some specific relations that might hold between experiences (and between events in general), and which are reasonable candidates for being a kind of part/whole relation; we won't worry too much about whether they are "really" instances of the part/whole relation. By "reasonable candidate" I mean that each relation gives us a partial order on experiences and is a relation of metaphysical containment or inclusion of some kind. It's not clear that these relations satisfy supplementation, or an analogue of this condition; this is discussed below in section 3. If they

don't, this makes them less paradigmatically mereological than, say, part/whole on spatial regions.

As mentioned, I'm assuming that experiences are events involving objects instantiating experiential properties, perhaps relative to times, so experience parthood is a species of event parthood. The object involved here—which I'll call the "logical subject" of the experience—may be coincident with a whole human body, but we should leave open the possibility that it is really a spatiotemporal part of a body, such as a brain, or a temporal part of a body or brain. An austere view is that an experience is fully individuated by a logical subject, experiential property (its "type"), and time. As we will see, there may be motivations for a richer conception. For one thing, it is worth noting right away that we tend to think of events as individuated in terms of the determinate way in which the event type is realized. For example, if I'm actually dancing at *t* by waltzing, but had I instead been dancing at *t* by belly dancing, then intuitively this would have been a different dancing despite involving the same subject, property, and time. Similarly, if I'm experiencing blue by looking at the sea at *t*, then had I instead been experiencing blue at *t* by staring at the sky, this would have been a different experience. We might helpfully distinguish *states* from *events*; states are fully individuated by a (perhaps determinable) property, subject, and time. Thus, in our first example, we have the same state of my dancing in both scenarios but different dancings.

On the *type-entailment approach* to event parthood, one event is part of another if and only if (1) they have the same subject and time (2) the type of one event entails the type of the other. Thus event type-entailment is directly defined in terms of *property* entailment. For example, perhaps it is necessary that if a person is dancing then they are moving; so the event of the person dancing at time *t* type-entails the event of them moving at time *t*. Bayne and Chalmers (2003) consider type-entailment as an account of experiential part/whole, although they do not ultimately endorse it. The view is fairly plausible because in general a complex experiential property does entail the simpler experiential properties involved in the parts of an experience. For example, the property of experiencing a gray cylindrical figure entails the property of experiencing a gray figure and the property of experiencing a cylindrical figure.

Note that entailments between properties are grounded in different ways; for example, the relation between a determinate property and its determinates is different from the relation between a complex property like a conjunctive property and any entailed properties out of which it is

defined. On the type-entailment approach, we should ask how entailments between phenomenal properties are grounded.

On the main alternative to type-entailment I want to consider, the *realization approach*, the mereological relations between experiences (or events in general) are defined in terms of mereological relations between their physical realizations. So in the cylinder example, the idea is that the physical realization of your experience of each cylinder is a part of the physical realization of your experience of the whole image, and so in a derivative sense one experience is part of the other.

A realization of an event (as I will understand it here) is a minimal set (or mereological sum) of events of a certain type (e.g., of a neural type) that are sufficient for it to exist.[3] A realizer is one of these events. So for example, a neural realization of an experience is a minimal set or sum of neural events that are sufficient for the experience to exist, and a neural firing might be one of its realizers. There are tricky questions about how exactly to understand realization (for discussion see Gillett, 2002; Kim, 2000; and Shoemaker, 2007), which I do not have space to discuss in full detail here. I will assume that there is a well-defined notion in this ballpark that can play the role I have in mind for it here. I will mostly deal here with what are often referred to as *total* realizations—realizations that are strictly sufficient for what is realized to exist. But the notion of *core* realization will be relevant too, a core realizer of an experience being the part of the total realizer that is relevant to fixing the specific type of the experience rather than the part of the total realizer that puts in place more generic conditions required for an experience to exist. For example, activity in V4 might determine the specific kind of color experience you are having, even though activity outside V4 (e.g., activity along thalamocortical loops) might also be required as a background enabling condition for the experience to exist.

If the realizers of the events we are interested in have a well-defined mereological structure, then we can define derivative mereological relations on the events themselves. It is fairly plausible that the realizers of experiences in particular *do* have a straightforward mereological structure. Suppose in a given case that there are a certain number of independent[4] neural events that we can take as "atomic" events. These will each involve a given neuron or neurons having a certain property, or a group of neurons related in a certain way. A fusion of any two of these atomic events will exist in a straightforward sense. For example, if event 1 involves N1 being F at $t1$, and event 2 involves N2 being G at $t2$, then their fusion is the event of

N1 being F at *t*1 and N2 being G at *t*2.[5] Hence we can think of our basic neural events as generating a mereological hierarchy of complex neural events. Assuming we have an adequate base of atomic events, the total realization of an experience will be a mereological sum of these atomic events. The realization of one experience can therefore be a part of the realization of another in a straightforward sense.

There are some important differences between the approaches. Note that in general, there is no requirement that if one event is part of another event then they have the same logical subject or happen at the same time, as is required on the type-entailment approach. For example, Bhoomika's dancing might be part of a performance by a dance troupe, despite their having different logical subjects and her dancing happening at a different time from other parts of the performance. The realization approach can easily allow for this, because the realization of her dancing might be a part of the realization of the whole performance. In my view, it is not unreasonable to think that experiences might turn out to be like this. For example, it might be that visual experiences are not individuated in terms of the whole body but rather in terms of brain areas that include the areas where visual processing happens, whereas auditory experiences are individuated in terms of areas that do not include these visual areas; so different parts of a total experience might have different logical subjects (see Lee, manuscript, for discussion). Or it might be possible for a single unified field of experience to have distinct experiences as proper temporal parts (e.g., Dainton, 2006, defends such a view). If this is correct, it is a reason why the realization approach is preferable.

Notice that it is possible to instantiate a determinable phenomenal property—such as experiencing the color gray—more than once at a single time; indeed, this is exactly what happens when you look at the gray cylinders. What distinguishes these experiences (call this the "duplication problem")? A plausible answer is that their determinable types are realized in different determinate ways; for example, each experience of gray involves a more determinate experience of a particular kind of gray cylinder at a specific location. A related problem: the type-entailment theorist won't want to say that each gray experience is part of the other, despite their types being mutually entailing. This problem can also be solved by taking determinable experiences to be individuated by their determinate realizers and part-whole structure to hold primarily between determinate parts (Bayne & Chalmers [2003] make the same point). This requires that there is a way of distinguishing the determinate parts of an experience (an issue I return to below in sec. 4) and also that determinate phenomenal types

are nonrepeatable within a subject's experience at a time. The appeal to realization also brings the type-entailment account close to the realization account; the main difference is that here we are appealing to realization *at the experiential level* (by a determinate phenomenal property instantiated by the same subject), whereas the realization approach descends down to the subpersonal level; it would distinguish the different bluish experiences in terms of their distinct subpersonal realizations.

Both the type-entailment approach and the realization approach are reductive accounts of the part-whole relation on experience, specifying the conditions under which it holds. Why think that such an account is possible? Perhaps the part-whole relation between experiences is *primitive*. Certainly, in the case of material objects and their parts, it is quite plausible that the relevant part/whole relation can't be cashed out in other terms. Furthermore, in our account of realization, we appealed to the notion of a *fusion* of two events, a notion that probably can't be defined in other terms, suggesting that at least sometimes part/whole on events *is* primitive. Two points to make about this: first, part/whole could be a primitive *determinable* relation, with certain determinate relations as instances—in this sense, the two accounts considered might not be incompatible with primitivism. Second, even if the part/whole relation between experiences isn't identical with either of the relations considered, it's plausible that it is at least *co-extensive* with one or another of them in the cases we are interested in, and that thinking in terms of these more substantive conditions is more informative. This suggests that we can reasonably set aside the primitivist view for present purposes.

Having put on the table some ideas about how to think about experiential parthood, let's now move to considering the main distinction that I want to discuss here between holism and atomism.

3 Priority: Holism and Atomism

Phenomenal holism is the view that the experiential parts of an experience exist only in virtue of the whole existing. Each part of the whole will be a phenomenal property instantiation; for the holist, this part exists only in virtue of the instantiation of a total experiential property. Intuitively, it is an aspect of, or abstraction from, the whole.

One version of holism is a property dualist version. The property dualist believes that some phenomenal properties are fundamental properties, whose occurrences are basic in the way that fundamental physical occurrences are basic. On the holistic version of this view, it is total phenomenal

properties that are fundamental. This view can be thought of as analogous to a quantum holist view, on which the fundamental physical properties belong to the whole universe and more local states of affairs obtain in virtue of the global state of the world.

Another version of holism is a physicalist version on which all phenomenal occurrences happen in virtue of physical occurrences. By way of analogy, consider the *center of mass* of all the neurons currently firing in your head. This is a three-dimensional quantity whose components depend holistically on all the neural activity happening in your head; a total phenomenal property could be a higher dimensional physical/functional property whose components depend holistically on a global pattern of neural firing in an analogous way.

Notice that on either version of holism, the "field" of experience and its phenomenal properties are quite different things from a spatial region and its physical properties; the "phenomenal field" is not itself a concrete particular, but rather involves a concrete particular such as a body, brain, or brain region having a certain complex, structured *property*. Therefore, "modifications" of the phenomenal field, especially modifications that change over time, are not analogous to the changing physical properties of a spatial region such as a changing magnetic field. Rather, they are the ways in which the total experiential property associated with a concrete object (like a brain) are changing over time.

Holism can be interpreted as entailing a *token modal dependence* thesis: each subexperience of a total experience necessarily only exists if the whole exists.[6] This token-level modal claim shouldn't be confused with a stronger claim of holistic *type*-necessitation, to the effect that an experience part of a certain type necessarily only exists in a total experience of a certain determinate type (a strong version of a kind of "gestalt unity." Bennett and Hill [this volume] discuss this type-level holism). The (token) holist could say: *this* intense toe pain owes its existence to the total experience I'm having right now, even if another intense toe pain of the same type could occur as part of a different kind of total experience. Compare how *this* instance of blueness may be constituted by, and so modally dependent on, an instance of royal blue, even though it is perfectly possible for a surface to be blue without being royal blue.

Suppose we assume that the basic neural events out of which we construct the total realization of an experience are modally independent: they could have existed without each other. Then, given that on holism the parts of an experience are mere aspects of the whole and so depend for their existence on the whole, it follows that each part of a total experience

has the same total realization (see endnote for more discussion).[7] Call this "realization holism."

Realization Holism: Each part of a total experience has the same total neural realization.

The fact that, arguably, the physicalist holist is committed to the parts of an experience having the same total realization should give us reason to question how likely it is to be true. I think that an argument of the following form is often very plausible:

(1) Parts E1 and E2 of total experience T have distinct core realizers.
(2) If E1 and E2 have distinct core realizers, then they have distinct total realizers.
(3) Therefore, E1 and E2 are prior to T.
(4) Therefore, holism is false.

Perhaps some holists will reject premise one in every instance, but it is often very plausible. Take an experience with auditory and visual components. We know that these experiences are associated with different brain areas, and this is fairly good evidence that they are at least core-realized in different places (I think this remains true, even once we acknowledge the role of intermodal feedback in perceptual processing: at most, this shows that in some cases the experience of a given feature is core-realized across different processing streams [see, e.g., Shimojo & Shams, 2001; see also Brogaard, Marlow, & Rice, this volume]). It is also plausible that a total experience is partially realized by one or more *structured representations* in the brain and that different parts of this representational structure are relevant to different parts of the experience, supporting distinctness of core realization in a different way. Distinct core realization doesn't *imply* distinct total realization, but on the other hand, even acknowledging the role that intermodal feedback sometimes plays, it remains plausible that typically at least *some* of what is happening in auditory cortex, which is relevant to your auditory experiences, is not relevant to the existence of your visual experiences, so that they have distinct total realizations. I take this to show that we have good reason for taking seriously alternatives to the holistic picture.

What anti-holist views are available? The weakest is simply a denial of holism. This is consistent with neither whole nor parts having priority, a view I will set aside here.[8] I will consider only "atomist" views, on which whole experiences have experiential parts that don't depend on the whole. These atomist views are distinguished by how strongly atomistic they are, in ways that I will categorize in what follows. To begin with, consider

the strong view that parts of experiences are *always* prior to wholes. This requires that experiences have basic parts that don't have any experiential parts (or perhaps infinitely descending decomposition). What would such basic parts be? One candidate might be experiences corresponding to primitive feature representations such as experiences of edges, colors, or textures (e.g., Zeki [2003] believes in such "micro-experiences"). But surely a view in the spirit of atomism need not be quite so strongly atomistic. It could be that there are *local holisms* in experience—for example, the experiences of the shape, color, and texture of a surface could be holistically interdependent (e.g., have the same total physical realization)—even though typically a total experience does have at least some independent parts (i.e., parts whose existence and nature do not depend on other experiences). I suggest we construe atomism as consistent with local holism and define it as the view that total experiences typically have *some* parts whose existence and nature are not grounded in the whole. We might call this "weak atomism" to emphasize that some of the connotations of "atom" are misleading here; I will omit this qualification in what follows.

On the atomist view, the total realization of a typical experience will itself have proper parts that themselves realize experiences; I'll call these "subrealizations." As above, I'll assume that two distinct subrealizations are modally independent (that is, one could exist without the other—unless one is a part of the other; note that overlap is consistent with modal independence). It follows that the experiences associated with these distinct subrealizations will be modally independent, unlike the parts of an experience on the holist view. So the construal of atomism in terms of experiential parts having metaphysical priority over the whole is very closely related to a construal in terms of modally independent experiential parts realized by distinct, modally independent subrealizations.[9]

Some further varieties of atomism I want to highlight arise from the distinction between the *content* of a conscious experience and the *fact that it is conscious*. This distinction leads naturally to a distinction between *consciousness atomism* and *content atomism*. According to some theorists, the character or content of an experience and the fact that it is conscious at all are separable features of an event in the sense that the consciousness of the event is in some way *detachable*. For example, functional theories of consciousness such as higher-order theories (HOT) and access theories seem to imply a kind of detachability because being metacognized or accessed by working memory is a contingent property of a representation.[10] If consciousness and content/character are separable in this way, then it is possible to take holistic/atomistic attitudes about them separately. In particular,

you could combine content atomism and consciousness holism: the content of an experience could be built up from below, even if consciousness is a property that belongs holistically to a total experience and only derivatively to its parts.

Note that by adopting a content atomist view, physicalist holists can make their commitment to realization holism more plausible. Content atomism allows the holist to say that parts of a total experience have distinct *core* realizations even though they have the same total realization.[11] It might also allow them to explain such features of total experiences as their variable dimensionality in a systematic way (see below). If they go this route however, then they have a version of holism that doesn't have the radical methodological consequences for consciousness research suggested by some holists like Searle. As mentioned above, Searle (2005) suggests that trying to explain consciousness from the bottom up by locating neural correlates of different aspects of experience is wrongheaded because consciousness has a holistic structure that requires theorizing about it as a whole field. However, if content atomism is true, we can at least hope to explain the *content* of experience in a piecemeal way, perhaps in part by locating local neural correlates, even if we need to think about consciousness itself more holistically.

I should stress that consciousness holists are still committed to the dubious claim that different parts of a total experience always have the same total realization—a claim that, if anything, becomes more dubious once we accept content atomism. I expect that much of the intuitive appeal of consciousness holism is shared by a different, weaker view that I call the *top-down view*, and which I will argue atomists can subscribe to. This is the topic of section 5.

What would be some concrete examples of holistic and atomistic views of consciousness? Most existing theories of consciousness either imply atomism or have atomistic versions. One big divide in theories is between "localist" views such as Block's (2007) and Lamme's (2003) on which perceptual experiences are wholly realized in areas concerned with processing perceptual information[12] such as the visual areas in the occipital lobe, versus "centralist" views on which areas involved in more "central" forms of processing, such as working memory, are involved in the realization of experience. Examples of centralist views are Baars' (1997) global workspace view, Dehaene and Naccache's (2001) neuronal implementation of this view, Prinz's (2012) attentional view, and HOT theories such as those of Armstrong (1968), Rosenthal (1997), and Carruthers (2000). Localist views imply that experiences associated with different modalities have separate

total realizers and therefore that atomism is true. Centralist views are also consistent with atomism: my visual and auditory experiences might have different total realizations even if these realizations *overlap* because the experiences exploit some of the same central resources (see sec. 6).

Concrete illustrations of the holistic approach are harder to come by. We can abstractly conceive of content holism being true: even if the contents of, say, a visual state are fairly well *correlated* with local visual activity, it could be that the correct content-determining principles assign content holistically to a total conscious state (imagine, for example, that they require consistency across and within modalities). Even supposing that content holism fails, we can envisage versions of consciousness holism. For example, there will be versions of centralist views on which the availability of a representation to central systems is determined holistically; for example, the relevant kind of central availability might apply primarily to a coalition of representations that are linked in the right way and only derivatively to each member of the coalition (Van Gulick's [2004] "Higher-Order Global States" view might fit into this category).[13] Thus the total realization of, say, a visual experience might include the realization the whole coalition it is part of, including any *non*visual representations that are members. In this way, what is happening in the auditory system might, after all, be relevant to the existence of your visual experiences.

Even though existing theories tend to be weakly atomistic, we shouldn't rush to reject holism; we probably do not know enough about consciousness to know for sure what mereological form it takes. Having said that, in what follows I want to shine some light on what does and does not follow from rejecting holism. As I said, my main point in this paper is that this does not require accepting an (perhaps) implausible "building- block" model. To that end, I consider in the following sections various theses that are suggested by the building-block metaphor: supplementation, constructivism, disjointness, and unity externalism, but which I will argue are not implied by atomism as I have defined it here.

4 Decomposition and Supplementation

How does a total experience decompose into parts? I want to consider the different answers we might get from holists and atomists. To that end, we need to first ask: how should holists and atomists think about part-whole relations on experiences? Because holists think that the parts of an experience have the *same* physical realization, they are likely to think in terms of type-entailment: a total phenomenal type will entail many less-specific

phenomenal types, which correspond to the parts of the experience. Furthermore, the holist should take the relevant entailment relation to be akin to the determinate/determinable relation, not a definitionally grounded entailment relation (assuming that a complex defined property is *derivative* from the simpler properties it is defined out of).

The atomist, by contrast, thinks that (at least some of) the parts of an experience have distinct physical realizations and therefore could think of experiential parthood in terms of mereological relations between these realizations. The same point applies to the mereological relations between the *core* realizers of experiences, which might be appealed to by a content atomist (this suggests that a consciousness holist who is a content atomist could also adopt a version of the realization approach). If the atomist does nonetheless think of parthood in terms of type-necessitation, it is open to her to hold that parts of a total experience can be definitional parts rather than determinable parts and so more basic components of that experience in a very clear sense.

As mentioned, our intuitive notion of parthood includes a commitment to supplementation. A proper part is accompanied by another disjoint proper part, and perhaps the stronger condition of complementation: there is a disjoint accompanying part that makes up the "difference." For present purposes, a helpful intuitive way to think of supplementation is as the view that every proper part of a whole is one of a set of disjoint parts that sum to the whole. If we assume unrestricted fusion, this is the same as complementation (in what follows, for simplicity, I will ignore the distinction between supplementation and complementation). The question I want to address here is: do total experiences break down into disjoint parts in an analogous way? Intuitively, it is only if some such condition holds that the "building-block" metaphor seems apt, so it is important for our purposes to consider the senses in which holists and atomists may or may not be committed to supplementation.

Consider first a version of holism on which the parts of a total experience correspond to those phenomenal properties that are type-entailed by the total phenomenal property. If we include determinable ways of experiencing as "parts" here, then we won't get a complementary structure: e.g., there is no experience that is the "difference" between determinably seeing something as blue, and determinately seeing it as royal blue.[14] Nonetheless, there is, intuitively, a sense in which such a holistic experience has *determinate* experiences as parts that might break down in a complementary way (furthermore, to solve the duplication problem [see above], the holist *needs* a way of distinguishing such determinate parts). One way to

cash this out is as follows: Suppose that a total phenomenal property is a high-dimensional quantity with a certain number of independent degrees of freedom. Much as a quantity like three-dimensional location has three independent components, we could think of the determinate parts of an experience as the ways in which it is determined along each phenomenal degree of freedom.[15] For example, suppose that a total experience consists of experiencing different colors at different locations and nothing more (an implausible "pixel map" model); the degrees of freedom might then be the ways of specifying the color at each location. Note that this satisfies supplementation in the sense that the total phenomenal property is the *conjunction* of the ways it is independently determined along each phenomenal dimension (although the whole is prior to each such part (compare: a 3-D location and its components)); we can think of the instantiation of each conjunct as a disjoint experiential part. By contrast, experiencing royal blue is not the conjunction of experiencing blue and some other independent experiential property.

One problem with this way of decomposing an experience is that these degrees of freedom, assuming they exist, are probably too fine-grained to correspond to our intuitive idea of a determinate experiential part. For example, color experiences have at least three separate degrees of freedom: brightness, saturation, and hue. But we would not intuitively say that we have separate experiences of these different aspects of a color. Perhaps determinate experiential parts correspond to *groups* of determination dimensions—but if so, it is unclear what the principle of grouping is. In response, the holist can point out that since they are already committed to denying that total experiences have independently existing experiences as parts, it is not a huge cost to them to concede that the determinate "parts" of a holistic conscious field do not correspond very precisely with our pretheoretical notion of an "experience".

Another issue is that total experiences have variable complexity, which presumably would correspond to variable dimensionality. One might hope for an explanation of this variable complexity at a deeper level in terms of the way different total experiences can be built from more or less independent components; but this seems to require a way of picking out these independent components other than *as* determination dimensions of the whole (this might be thought of as a general objection to holism). Relatedly, unlike fundamental physical quantities, it is far from clear that total phenomenal properties can be neatly broken down into dimensions that are completely independent; for example, the hue, brightness, and saturation of color are *not* completely independent.[16] Again, we might hope for

an explanation of this at a deeper level. In response, perhaps a physicalist holist can explain these features at the level of subpersonal realization (for example, they might appeal to representational structure, perhaps by being a content atomist); the property dualist holist, on the other hand, seems to have no choice but to take the variable structure and interdimensional dependence of total experiences as basic and inexplicable.[17]

The main point to emphasize here, however, is that for the holist, these determinate dimensional parts, even if they involve independent variables, are not independently *existing* experiences (they all depend on the whole) and therefore not disjoint building blocks in the sense that the holist was concerned to reject. An atomist could also think that total phenomenal properties are separable into independent dimensions, but she is also unlikely to think that every such dimension corresponds to an independently existing experience (e.g., your experiences of the hue and brightness of a color are not independent existents). This raises the question of whether the atomist is committed to an experience having determinate complementary parts that are independent or "disjoint" in a stronger sense.

Consider the set of subrealizations associated with the realization of a total experience. For each such subrealization there will be a maximally determinate experiential property that it realizes, which entails all the other experiential properties it realizes. You might think this gives the atomist a nice way of breaking up a total experience into independently existing, fully determinate parts. However, we need to be careful here. We haven't yet said that the atomist is committed to such experiential parts having complements; if they don't, then we can't "break down" a total experience in this way.

In what sense might these fully determinate, independent experiential parts have complements? We can certainly consider their "realization complements": the complement of their total realization within the total realization of the associated total experience. However, it is not clear that a realization complement will itself realize any experiences, and if it doesn't, then it will not give us an *experiential* complement; for example, it might be that all experiential parts of a total experience have overlapping realizers (see sec. 6). There is another more liberal definition of "experiential complement," however, that doesn't require experiential complements to have disjoint realizers. Let's suppose that T has independent experiential parts E1 and E2. They are *experientially disjoint* if they don't share any experiential parts (i.e., the intersection of their realizations doesn't realize any experiences). One can be thought of as the complement of the other if (1) they are experientially disjoint, and (2) they "cover" the total experience:

that is, either (2a) the fusion of the realizations of E1 and E2 is coincident with the realization of T, or (2b) T is constituted by E1 standing in the phenomenal unity relation to E2. The reason for the disjunction between (2a) and (2b) has to do with the issues discussed below, of whether constructivism and unity externalism are true. For now, it is probably helpful to think of complements as satisfying (2a). I'll say more to explain the difference between (2a) and (2b) below.

Unfortunately, there is no guarantee that an experience will have a complement, even in this more liberal sense. For example, a total experience of a table against a white background might have the experience of the table as an independent proper experiential part but have no other independent parts that are experientially disjoint from the table experience (e.g., the experience of the background might not be an independently existing experience). Thus it is far from clear that merely granting that a total experience has some independent experiential parts implies that it can be fully broken down into a set of experientially disjoint parts.

Let's spell this out in more detail. Suppose your current experience does break down into independent parts. What are they? We can distinguish a *bitty* view and a *chunky* view. On the bitty view, the parts involve the experiencing of simple features such as colors, edges, and textures as placed in space in certain ways or as attributed to certain objects or surfaces and perhaps also experiences of the spatiotemporal relations between objects and events. "Simplicity" might be thought of in phenomenological terms, or, perhaps more promisingly, we might postulate that experience is underwritten by a structured sentence-like representation, and the features correspond to primitive predicates in a kind of "perceptual language" used, e.g., by the visual system. There are two problems with this that I want to mention. First, even if experience is underwritten by structured representations of some kind, it is unclear that they will have a canonical decomposition into basic parts involving simple predications; as Fodor (2010, ch. 6) points out, there are nonsentential forms of representation—such as imagistic representations—that may resist such decomposition even if their content is a function of their structure. As far as I know, it is an open question whether the representations underwriting conscious experiences (assuming there are such things) are like this. Second, even if we grant such a decomposition involving basic predicates is possible, it is dubious that such basic feature attribution underwrites independent experiences. In particular, even though it may not generally be true that features have to be bound to objects to be experienced,[18] when binding *does* occur, it is plausible that we have a kind of experience of a feature that could not exist outside of

binding, and whose existence depends on the other features that are experienced as bound to the object (think, e.g., about the interplay between color, shape, and texture). It is therefore implausible to think that feature experiences could be independent building blocks.

On the alternative chunky approach, the smallest independent experiences are larger parts of experience such as experiences of whole objects having multiple features (as such, there will be "local holisms," perhaps for quite large parts of experience). One problem with this approach is arbitrariness: it is unclear that there is the right kind of natural joint in experience at an intermediate level. Why an experience of a whole object rather than an experience of a surface, part of the object, or the relations between multiple objects? Despite this problem, it remains plausible that total experiences do have at least *some* independent parts, so presumably there are *some* such joints, even if they don't cleanly divide experience into complementary parts. One possibility is that the smallest independent parts are modality-specific fields of experience. However, it may be that the experiences of some features are realized across different modality-specific processing streams (as mentioned above, this is one possible interpretation of what happens when there is intermodal feedback concerning a single feature). If so, then the real joints are more likely to be at the level of distinct object representations, or representations of distinct spatiotemporal regions, whose features may be interrogated in a multimodal way.

Supposing there are such independent parts at a "chunky" level, it is unclear whether they have complements, for reasons already gestured at above. Take experiences of relations (such as spatial distance) between objects or regions. Even if the object experiences involved are independent experiential chunks, it is unclear whether my experience of the distance between the two objects could exist on its own without any of my experiences of the intrinsic features of the objects. The same point applies to certain "gestalt" experiences such as perceiving the cylinders in figure 13.1 as forming rows rather than columns; plausibly, your experience of a row formation requires that you also have experience of at least some intrinsic features of the cylinders, such as their shapes; so your "rowish" experience is not an independent "chunk." A similar point was made above concerning the backgrounds of objects, and we could make the point about temporal relations between events and experiences that attribute high-level features based on lower-level features that are perceived (e.g., my experiences of the low-level features of a face might be an independent experience. If in addition I experience it *as* a face, this experience may not be independent of the low-level experiences).

To sum up: It is unclear whether experiences have bitty parts, and even if they do, it is implausible that these parts are independent experiences. If the smallest independent parts are at the chunky level, then it becomes unclear whether they have experiential complements. So if we are atomists, we have good reason to be skeptical about supplementation.[19] This is the first way in which atomists need not subscribe to a "building-block" picture.

5 Constructivism and the Top-Down Approach

The points in section 4 about supplementation relate closely to another issue concerning the proper interpretation of atomism. I mentioned earlier that we might think that there is a "unity relation" that holds between experiences, such that when they are unified, they are components of a more complex experience. Furthermore, many philosophers assume that specifying the nature of this relation is an independent task for a theory of consciousness. We might think that an atomist in particular is committed to saying what this relation is, as it is needed as a kind of binding agent for fusing together atomic experiences into wholes. I want to argue that atomists needn't picture things like this.

In the literature on the unity of consciousness, you can find a menagerie of different accounts of what the unity relation is. One large class of accounts treat unity in functional terms as a kind of integration between the contents of experiences (Shoemaker, 2003; Prinz, 2012); for example, experiences might be unified if their contents are jointly available to a single working memory system, or jointly represented by a higher-order monitoring system (Rosenthal, 2003). A related approach is the conjunctive representational approach, on which experiences are unified just if they contribute their contents as conjuncts to the complex content of an overarching experience (Hurley, 1994; Tye, 2003). On a subject-based account, unified experiences are those that belong to a single subject of experience (McDowell, 1997). Or on a primitivist view, unity is a primitive relation that can't be analyzed in other terms (Dainton, 2006). There are other approaches as well.

The idea that we need to specify some such relation would certainly be correct if a constructivist view of complex experiences is correct. The view has two components. First, there must be a basic supply of independently existing experiences that we can treat as "simple." These simple experiences might be bitty, and so "simple" in an intuitive sense, but they could also be chunky (see above). Second, a "complex" experience is an event consisting in the obtaining of the (independently specified) unity relation between a

set of simple experiences. For example, they might be simple experiences that are functionally integrated in the right way. Thus complex experiences are really "experiences" in a derivative sense: they are complexes that inherit their status as experiences from the simples out of which they are constructed.

One noteworthy version of constructivism is unity pluralism (see Bennett & Hill, this volume;[20] and Hill, 1991, ch. 10). Unity pluralists think that it is a mistake to think that there is a single, privileged unity relation that holds between experiences. Rather, there are myriad types of integration and unification that experiences can enjoy—such as those mentioned above—and the best we can expect from a theory of the "unity of consciousness" is a theory of each of these different relations. I think that once we accept the constructivist view, this is a plausible way to go: why think that one way of forming complexes from simple experiences has some privileged status? The theory also has the merit of explaining why we find it hard to say in certain cases whether a complex state is an "experience:" it might be unified in some ways, but not in others.

I believe that there are at least some uses of the term "experience" for which a kind of constructivist approach is correct. Consider the experience of watching an entire movie. There is a sense of "experience" on which this is a single experience. However, arguably, in another sense this is just a temporally extended series of experiences that stand in the intimate relation of being enjoyed by a single subject during a certain period of time. We might take a constructivist approach, treating temporally local experiences as "simples" and allowing that whole-movie experiences can be "experiences" constructed from these simples in a derivative sense. The general constructivist thinks that an analogous approach is correct for complex experiences at individual moments.

Constructivism requires supplementation (with some qualification).[21] If supplementation fails, then we won't have the requisite supply of independent building- block experiences out of which we can construct more complex experiences. I think this makes constructivism implausible for the same reasons. Consider again the gestalt experience as of cylinders arranged in rows rather than columns (fig. 13.1). Even if the experiences of cylinders are independently existing parts of the experience, we can't think of the gestalt experience as obtained by gluing an independently existing "row-ish" experience to the individual cylinder experiences.[22] In this sense, it is not constructed from below.

Fortunately, there is an alternative approach to unity and composition that is independently attractive, even if supplementation does obtain. On

what I'll call the "top-down" approach, we can formulate a theory of what makes something an experience that applies *directly* to the range of complex experiences we are interested in so their status as experiences does not derive from their parts. Furthermore, rather than thinking that "complex experience" is defined in terms of an independently specified unity relation, we now take "unity" to be defined in terms of "conscious state": two states are "unified" just if they are components of a state that is conscious. It's helpful to think of this in terms of the unity biconditional (see Bayne & Chalmers, 2003):

Unity Biconditional: Experiences E1 and E2 are phenomenally unified if and only if they are components of some complex experience E3.

The constructivist takes the LHS to ground the RHS, whereas the top-down theorist thinks about things in a more "top-down" way, taking the RHS to have priority. What's nice about this approach is that once we have a theory of what makes a state conscious and what it is for one experience to be part of another, we get a theory of unity for free—so, saying what unity is turns out *not* to be a separate task for a theory of consciousness. This is the approach advocated by Bayne and Chalmers (2003) and Bayne (2010). Bayne calls it the "mereological" view of unity.

To see an example of the top-down approach in action, consider views on which phenomenal consciousness is really a functional property such as access consciousness (the property of being available to postperceptual consuming systems involved in motor planning, memory formation, and reasoning) or higher-order consciousness (the property of being targeted by a higher-order representation). On the top-down view, each functional account of consciousness delivers a functional account of what phenomenal unity is—for example, if phenomenal consciousness is access-consciousness, then phenomenal unity is access-unity (i.e., two states are phenomenally unified just if they are parts of a state that is access-conscious, i.e., just if they are jointly accessible to consuming systems).

Another important example is the view of unity generated by a representational view of consciousness on which *all* experiences—including complex experiences—are a species of representational states that present the world as being a certain way: i.e., they have a particular propositional content (which may be further taken to determine the qualitative character of the experience). The representational theory combined with the mereological approach naturally generates the view that phenomenal unity is *conjunctive* unity: experiences are unified when they contribute their contents as conjuncts to the complex content of an experience that subsumes them.

This view is noteworthy because it promises to reduce experiential mereological structure to a certain kind of *representational* structure, conjunctive structure (although for some objections to it, see Bayne, 2010, ch. 3).

In addition to avoiding the potential problems for supplementation mentioned above, one advantage of this approach is that it gives a *unitary* account of what consciousness is for both complexes and more simple experiences. The constructivist approach says, in effect, that something is an experience if it is either a simple experience *or* is a complex constructed in one or more ways from simple experiences. The top-down approach avoids being disjunctive in this way, which would be a nice feature of a theory of consciousness if we can get it.

This plays out in an interesting way in a recent discussion of unity in Prinz (2012, ch. 7). Prinz's theory is that representations are conscious when they are attended (which for him means being accessible to working memory); he then proposes that attended representations are unified when they are realized by neural populations that have phase-locked firing patterns—unity is "neural resonance." This initially sounds like a constructivist approach: the "simple" experiences are the attended ones, and the "complex" experiences are collections of attended states realized by resonant neural populations. However, Prinz is careful to argue that his account is really more unitary because neural resonance can be understood as realizing a form of attention, "co-attention," that can apply to complex multimodal wholes. So this might after all be a top-down approach.

The above examples illustrate how the top-down "mereological" view of unity is not necessarily in tension with other views on unity in the literature. Rather, they can be seen as ways of implementing the approach. Clearly, whether the top-down approach can be made to work (that is, whether complex experiences directly satisfy the condition for being conscious) depends on what the correct theory is of what makes a state conscious, and I won't be able to address this big question here. The point I want to emphasize here is that *atomists need not be constructivists, but could instead adopt the top-down view.*

Consider, for example, the view that phenomenal consciousness is access consciousness and play it through the top-down account to get the view that phenomenal unity is access unity. On this view, what makes a complex experience an "experience" might be that it is access-conscious, not that it has access-conscious parts that stand in a separately specified unity relation. Furthermore, a complex access-conscious whole could have parts that have distinct total realizations and which involve distinct access-conscious representations (e.g., representations in visual and auditory areas) so that

the complex whole has distinct atomic experiential parts (this could be true even if what makes one part access-conscious is not independent from what makes the other part access-conscious; see below). This illustrates how atomism is perfectly compatible with a top-down approach. This is the second way in which the building-block metaphor may be a misleading way of picturing the atomic view.

Note that holists are automatically committed to the top-down view because they deny that total experiences have any independent parts. If atomists adopt the chunky view rather than the bitty view, then, in effect, they think there are at least some relatively simple experiences for which holism is locally true and for which the top-down approach is therefore also true. As mentioned earlier, most theorists will probably also accept that there are some "complex experiences"—such as experiences of whole movies—for which a restricted constructivist view is true. So there is a question for an atomist who thinks that the top-down view is *sometimes* true concerning which "complex experiences" it applies to and which complex experiences are really only "experiences" in a derivative sense. We should be open to the possibility that the concept of "consciousness" is indeterminate in such a way that it is not totally clear where exactly this line is drawn.

As I discussed at the beginning, thinking that the atomist must believe in supplementation and thinking that they must be constructivists are not the only ways of saddling the atomist with a kind of building-block conception of the structure of experience. Building blocks are *disjoint* objects, and they require *external* cement to hold them together. In the next section, I explain how there are plausible versions of atomism that deny that the parts of experiences are separate in the way bricks are and deny that any external glue is required for unification.

6 Atomism, Unity Overlap, and Internalism

Although atomists think that experiences typically have independent parts, we should be careful about what this independence amounts to. I defined it as existence independently of other experiences, which translates at the level of realization to the claim that two experiences are mutually independent if their realizations are distinct: that is, they are not coincident with each other, nor is one a proper part of the other. Holists think that the parts of a total experience always have completely coincident total realizations. We shouldn't think that the only alternative is independent parts that have completely *disjoint* realizations—what we might call "strong

independence." Experiences can have *overlapping* realizations even if they don't share any experiential parts,[23] that is, even if the intersection of their realizations does not realize any experiences. An atomist could hold that unified experiences are *weakly dependent* in the sense that they overlap.

Unity Overlap Thesis: Necessarily, if experiences E1 and E2 are unified, then they have overlapping realizations.

Such overlapping experiences would not be separate building blocks in quite the sense that bricks are separate parts of a house; the only completely independent building blocks of the total experience might be subexperiential events like individual neural firings. Nonetheless, they are parts of the total experience that are metaphysically prior to the whole and which can exist without each other. (More weakly, the building-block metaphor would be misleading even if the independent elements of a total experience *sometimes* overlap; that is, if "disjointness," the view that they are always disjoint, is false.)

We noted above that experiences of different types—e.g., visual and auditory—sometimes are core-realized in different areas. This suggests that it would not be plausible to require that the core realizers of unified subexperiences overlap, and that a more specific version of the unity overlap thesis might be correct, linking the consciousness-making parts of the total realizations of experiences (i.e., the complement of the core realization in the total realization).

Consciousness-Making Overlap Thesis: The consciousness-making parts of unified experiences overlap.

We can think of this as the intuitively attractive claim that if two experiences are unified, then what makes one conscious is not completely independent from what makes the other conscious.

Another important thesis in the same vein, and which the atomist could endorse, is unity internalism.

Unity Internalism: Unity is an internal relation between experiences.

For the internalist, the unity between two experiences is not some extra state of affairs in addition to the existence of the experiences themselves. Unity is an *internal* rather than an *external* relation between experiences: it holds in virtue of the intrinsic features of each experience, and therefore no substantial process of "binding" is required to make them unified. It is perhaps a little unclear what the "intrinsic" features of an experience are. I will assume that they include having a particular logical subject and

phenomenal type and being physically realized in a certain way. Note that it is plausible that these intrinsic features of an experience are also *essential* to it. If that's right, then we can take internalism to imply that once experiences exist, nothing more is required to make them unified (if they are unified at all). Unity supervenes on existence.

We can define *unity externalism* as the view that, except in certain degenerate cases, unity is an external relation. (We need the "degenerate cases" provision because one way for two experiences to be unified is for one experience to be part of the other experience. This is an internal relation.)

Internalism should not be confused with the view that phenomenal *types* are internally related through unity in the sense that a phenomenal type necessarily entails the others it is coinstantiated with, and vice versa. Dainton (2006) construes the idea that unity is an internal relation this way. This view is implausible—it implies, for example, that if a certain type of pain occurs in the presence of a thought about lunch, then that type of pain necessarily always occurs accompanied by a lunch thought. Maybe *some* phenomenal types necessitate each other in this way—I will not discuss this idea here; the general claim is surely too strong.

Phenomenal holism implies unity internalism because on holism, the subexperiences of a total experience only exist as unified parts of that total experience, so their unity supervenes on their existence. But atomists can embrace internalism as well. Consider, for example, the view that two experiences are unified if and only if they have overlapping realizations (I'm not saying this is a plausible view). Overlap is an internal relation between experiences (assuming the realization of an experience is an intrinsic property of it), so this would be a view on which unity is internal. Furthermore, it is clearly compatible with atomism: overlapping experiences need not be completely coincident (and therefore their existences need not be mutually dependent), as holism demands. Note that, as this example illustrates, internalism is perfectly compatible with two internally unified experiences being such that *each could have existed without the other*. It only requires that *if* they both exist, they are unified.

Let a fusion of two events, A and B, be an event whose realization is a fusion of the realizations of A and B (which we are already assuming is well defined). Given the assumption that a fusion of two events always exists,[24] internalism is equivalent to another mereological thesis about experiences—the unity fusion thesis.

Unity Fusion Thesis: Two experiences are unified if and only if there is an experience that is the fusion of these experiences.

If an experience fuses two experiences, then they are unified because they are parts of an experience (bear in mind that this is a nontrivial condition—not any two random experiences are parts of an *experience,* even if they fuse to form an event of some kind or other). Furthermore, if the fusing experience contains nothing other than the two experiential parts, nothing other than their bare existence was required for them to be unified (no external events were needed to bind them together). Thus unity fusion and internalism are equivalent.

A stronger thesis that connects these claims with the issue of supplementation raised above is the phenomenal decomposition thesis.

Phenomenal Decomposition Thesis: Each independent part of an experience is a member of a set of experientially disjoint experiences whose fusion is the whole experience.

This is equivalent to the conjunction of internalism and supplementation. If internalism fails, then a fusion of experiences won't in general be an experience. There will be extra events in the realization of a total experience needed to "link" the experiential parts. If supplementation fails, then a total experience won't necessarily be exhausted by its proper experiential parts, even though if internalism is true, then their fusion is itself an experience. If phenomenal decomposition *is* true, then a total experience is "built up" from its proper experiential parts in a very straightforward sense,[25] although we might be skeptical of this for the same reasons we might be skeptical of supplementation.

If externalism and supplementation obtain, then we get a somewhat different decomposition thesis. The existence of a total experience will consist in the proper parts of the experience existing and standing in substantive external unity relations to each other. These experiential parts won't fuse into a total experience because they need to be externally linked, much as a building is not a mere fusion of bricks because the bricks need to be cemented together.

What is the relationship between internalism and the unity overlap thesis? Arguably internalism is a stronger claim in the sense that it implies unity overlap, but there is no reverse implication. If two experiences do not overlap, then surely some external connection is required between them for them to be unified.[26] So internalism seems to imply unity overlap. The reason the reverse entailment—from phenomenal overlap to internalism—is less obvious is that it seems perfectly coherent for unity between experiences to require that they overlap, but *also* to require that certain external events relating the realizations of the experience to be in place as well. That

is, the unity relation might have both internal *and* external components. In practice, it's hard to imagine a view on which things work out this way, so I suspect that in the context of most theories of consciousness, internalism and unity overlap will either be jointly accepted or rejected. But nonetheless, unity overlap is arguably a weaker claim for this reason.

What, if anything, do existing theories of consciousness tell us about unity overlap and unity internalism? Localist theories like Block's and Lamme's suggest that a unified audio-visual experience can have separate components realized entirely in different perceptual systems, suggesting that these realizations are disjoint and that the relevant unity relation is external.[27] Centralist theories, on the other hand, are at least consistent with internalism and unity overlap. It might be that in order to be jointly accessible, unified experiences have to be partly realized in the same part of the central system (e.g., in the same working memory systems), and thereby have overlapping realizations—i.e., what makes one experience accessible is not independent of what makes the other accessible. It is also consistent with this that no external link between these realizations is required to make the experiences unified; the overlapping neural activity that contributes to both states being conscious/accessible might be sufficient on its own for the experiences to be unified.

Although centralism is consistent with unity overlap and unity internalism, it should be stressed that it is also consistent with them failing. To return to the example of Prinz's (2012) attentional theory on which consciousness is accessibility to working memory: Prinz emphasizes that working memory is not a single integrated system but a network of different systems and that representations may be conscious in virtue of being accessible to different components of the working memory network. This suggests that on Prinz's view, jointly conscious states need not have overlapping realizations, and their unity may be an external matter (which is consistent with Prinz's view that unity is neural resonance, an external relation between neural populations).

To sum up: Unlike bricks, the independent parts of an experience need not be disjoint, and contrary to the "cement" metaphor, unity need not be an external relation between experiences. It is also helpful to consider views like unity overlap and internalism (as well as the top-down view) because they show how a version of atomism can be developed that has some of the features of holism that I think some authors find attractive—e.g., the idea that the subexperiences of a total experience are not completely independent events.[28] Having said this, although unity overlap and unity internal-

ism are attractive views, existing theories of consciousness are mostly either neutral about them or in some cases rule them out.

7 Conclusion

I have tried to clarify the difference between atomism and holism and explain some different ways the atomic view might be developed. In particular, I have emphasized four different theses suggested by Searle's "building-block" metaphor that the atomist can reject: supplementation, constructivism, disjointness, and unity externalism. This raises the question of whether existing theories of consciousness support more moderate versions of atomism that reject one or more of these theses, and thereby have some of the attractive features of holism. Current theories are at best noncommittal on these issues, even theories where shared central resources such as working memory are implicated in making different kinds of experience conscious. We are still in the early days of understanding consciousness, however, so current theories may not indicate much about how more mature theories will lead us to think about these issues. I hope to have at least made clearer the different options available to an atomist and to have illustrated how the rather abstract concerns of most of the chapter play out in the context of more concrete theories of consciousness.

Acknowledgments

Thanks to Tim Bayne, John Campbell, and Dave Chalmers for helpful comments on an earlier draft.

Notes

1. At least, this is implied given the assumption that there are only finitely many parts in the relevant model and that sets of these parts exist. The complications that arise from dropping these assumptions are not relevant here.

2. This possibility creates an interesting methodological puzzle. We might think that the way to identify nonparadigmatic experiences (such as whole fields of awareness, if they count as experiences), is to figure out the nature of paradigmatic experiences and then identify the states that have the same real nature. However, if paradigmatic experiences are derivative aspects of nonparadigmatic experiences, then we won't be able to understand their nature without first identifying the nonparadigmatic experiences that they depend on.

3. Note that contrary to the approach of e.g., Kim (2000), I'm operating with a conception of realization on which realizer events need not involve the same logical subject as the property instantiation that is realized.

4. Independence is important. Without it, there is no guarantee that disjoint sets of atomic realizers do not realize some of the same phenomenal events. Independence at the neural level can be defined in terms of realization at a more basic level—neural events are independent if their fundamental, physical realizations are disjoint. This raises the question of how to understand independence at the fundamental physical level, and also whether there are any independent events at this level, and if so, whether they are the right kind of ingredients to build neural events from. I will not be able to pursue these important questions here.

5. If an object exists that is a fusion of N1 and N2, then we can think of this event as an instantiation by this two-neuron object of a structural property involving one part being F and the other G.

6. This assumes that it isn't contingent that a token experience is grounded in the whole. Although it is probably coherent to hold a grounding view on which these grounding relations are contingent, I think it is plausible to interpret a holist as believing that the grounding relations between parts of an experience hold necessarily—this certainly fits best with talk of parts as being mere aspects of, or modifications of, the whole field.

7. Suppose two parts of a total experience have different total realizations. Then, given the assumption that these realizations are constructed out of different sets of modally independent neural events, it follows that each realization could have existed without the other (assuming that one realization is not a part of the other), so each experiential part could have existed without the other, contrary to holism.

Two notes on this argument. First, I am assuming that on holism, it is not contingent that the existence of a token experience is grounded in the existence of the total experience of which it is a part. Second, there are views on which the different parts of a neural realization are not independent events. Consider a version of Schaffer's (2010) holism, on which local neural events are grounded in global neural events. On such a view, even if two parts of a total experience are grounded in completely disjoint and spatially localized neural events, they might nonetheless be grounded by the total experience, in that they are grounded in the total neural state of the brain. A holism on which parts of an experience can still have extremely localized total realizers is clearly not in the spirit of the views intended by actual holists, even if it technically counts as a form of holism, and so I will set it aside here. (Thanks to Dave Chalmers for discussion here.)

8. Shoemaker's (2003) view might be interpreted in this way. Whether there is a stable view intermediate between holism and atomism is an interesting question, but I won't address it here. I would at least point out that my definition of atomism is sufficiently weak that it is hard to find room for an intermediate view.

9. Note that this kind of modal independence is consistent with a given experience putting *constraints* on the kind of (determinate) experiences that can accompany it; the fact that E1 could have existed without E2 does not imply that E1 can exist on its own or that it can exist with any old experiences accompanying it.

10. I would suggest we understand detachability as the thesis that the experience has a realization part that fully determines its content/character and that could exist as part of an unconscious event.

11. This would also allow them to adopt a version of the realization approach to experience-parthood: they could say that parthood relations between experiences derive from parthood relations between their core realizers (obviously, they can't adopt this approach using total realizers).

12. I should note that the main concern of localists such as Block and Lamme is to deny that postperceptual frontal areas are involved in conscious perception; this is consistent with the total realization of experience not being limited to cortical perceptual areas but also involving connections with areas such as the thalamus, which may be involved in sustaining conscious awareness (see, e.g., Block, 2007, 482).

13. We need to be careful here to distinguish between a holism in the dynamics of selection for consciousness and a holism in the individuation of conscious states. By way of analogy, imagine an army of zombie soldiers swarming to enter into the fortress of consciousness. What determines which zombies get through the gates might be an extremely holistic process involving interactions across the whole group; nonetheless, once a solider makes it through the gate, what it is for it to be inside the fortress might be definable without reference to the other soldiers.

14. Note that there is a difference between explicitly experiencing the color as falling under the category "blue," and merely experiencing it as having a shade that is in fact a shade of blue. It is really the latter property that is the determinable of experiencing it as royal blue.

15. A complication here is that some multidimensional quantities, such as three-dimensional location, don't have a privileged coordinatization. Intuitively, total phenomenal properties are not in this category, but if this intuition is wrong, then even if they have a determinate number of independent dimensions, this wouldn't on its own determine a privileged decomposition into specific dimensions.

16. For example, some levels of saturation are not available for every hue (see, e.g., Palmer, 1999).

17. Perhaps dualist holists who take total experiences to be primitive relations to structured entities such as propositions can avoid this objection.

18. Possible examples of unbound feature experiences are color experience in a ganzfeld, motion experience in peripheral vision, and experiences had in situations where binding failures are prone to happen such as while viewing the stimuli that

prompt "illusory conjunctions" (Treisman, 1998). (I do not say it is obvious that any of these cases really do involve unbound feature experiences.)

19. A more complete discussion would also consider the mereological structure of *core* realizers; it might satisfy supplementation even if experiences do not have experiential complements in the sense we have been discussing. For example, if a form of content atomism is true on which a total experience is partly realized by a structured representation with various parts joined together through conjunction, these representational parts might have complements even if these complementary parts do not correspond to independently existing experiences.

20. There is a way in which Bennett and Hill's version of unity pluralism appears to differ from constructivism, as I am understanding it here: for them, not all complex states formed when experiences are unified (by one of the plurality of unity relations) should themselves count as "experiences." However, I think this is merely a verbal disagreement: my constructivist says that "experience" applies to these complex states, but only in an extended sense. Bennett and Hill say we shouldn't use the term "experience" to describe them. Clearly, nothing of theoretical importance turns on this.

21. Constructivism does not strictly imply supplementation. As I am understanding supplementation, an experience and its complement are experientially disjoint. I think we can make sense of a version of constructivism on which the "simple" building blocks out of which complexes are constructed are allowed to experientially overlap (which would require them to be somewhat "chunky"), even though they are capable of independent existence. I think this is an odd view, and I do not think it has significant advantages over a version of constructivism on which the building blocks are disjoint; but ideally it would be treated in detail as a separate case.

22. That is, this particular *token* "rowish" experience depends for its existence on the experiences of individual cylinders. That is not to say that one could not experience some different items as forming a row. In that case, we might have a rowish experience of the same *type*, but it would not be the same *token* rowish experience.

23. Note that overlapping in this sense requires more than having overlapping logical subjects; even if their subjects overlap, two events might not share any events as parts.

24. We need this assumption to rule out views on which the fusion of two experiences only exists under special circumstances such as their being related by an external unity relation. (Thanks to Dave Chalmers for pointing this out.)

25. Note that this is not equivalent to an even stronger thesis, to the effect that a total phenomenal property is equivalent to a conjunction of determinate phenomenal properties that characterize the parts of the experience. If holism is true, this

thesis may be correct if we can think of each of these phenomenal properties as a way of making determinate the total property along a determination dimension, but if atomism is true, then presumably a total phenomenal property is not a mere conjunction of more basic phenomenal properties.

26. Objection: Isn't belonging to the same subject sufficient for unity, and also an internal relation? But why think that experiences belonging to the same subject have to overlap? Response: Belonging to the same subject is only sufficient for unity if we individuate subjects in terms of unified fields (e.g., if subjects are just organisms, then there is no guarantee that their experiences are unified—e.g., an organism might have two separate brains). If two experiences belong to a single field-subject but do not overlap, then surely there is some external relation that makes them both unified and belong to the same subject. In other words, in the relevant sense, belonging to the same subject is only sufficient for unity if it is an *external* relation *or* the unified experiences overlap.

27. Although, as noted above in endnote 12, localists may in fact think that there are some neural structures such as thalamocortical connections that enable consciousness and which might be shared by different parts of a total experience. This version of the view is probably consistent with internalism and overlap.

28. Also, insofar as these theses illuminate the sense in which experiences have a mereological structure by focusing on mereological relations between the *realizations* of the experiences, they help motivate the idea that the mereology of experiences should be understood in terms of the mereological relations between their realizations, and furthermore that experiences are partly individuated in terms of their realizations.

References

Armstrong, D. (1968). *A materialist theory of the mind*. London: Routledge.

Baars, B. J. (1997). In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies*, *4*(4), 292–309.

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

Bayne, T., & Chalmers, D. (2003). What is the unity of consciousness? In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation* (pp. 23–58). Oxford: Oxford University Press.

Bennett, K. (2011). Construction area (no hard hat required). *Philosophical Studies*, *154*(1), 79–104.

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, *30*(5), 481–498.

Byrne, A. (2009). Experience and content. *Philosophical Quarterly, 59*(236), 429–451.

Carruthers, P. (2000). *Phenomenal consciousness: A naturalistic theory*. Cambridge: Cambridge University Press.

Dainton, B. (2006). *Stream of consciousness unity and continuity in conscious experience*. London: Routledge.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition, 79*, 1–37.

Fine, K. (2011). Towards a theory of part. *Journal of Philosophy, 107*(11), 559–589.

Fodor, J. A. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.

Fodor, J. A. (2010). *LOT 2: The language of thought revisited*. Oxford: Oxford University Press.

Gillett, C. (2002). The dimensions of realization: A critique of the standard view. *Analysis, 62*(4), 316–323.

Hill, C. S. (1991). *Sensations: A defense of type materialism*. Cambridge: Cambridge University Press.

Hurley, S. (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). London: British Academy.

Kim, J. (2000). *Mind in a physical world: An essay on the mind-body problem and mental causation*. Cambridge, MA: MIT Press.

Lamme, V. A. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences, 7*(1), 12–18.

Lee, G. (manuscript). Selfless experience.

McDowell, J. (1997). Reductionism and the first person. In J. Dancy (Ed.), *Reading Parfit* (pp. 230–250). Oxford: Oxford University Press.

Palmer, S. E. (1999). Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences, 22*(6), 923–943.

Prinz, J. (2012). *The conscious brain*. New York: Oxford University Press.

Rosenthal, D. M. (1997). A theory of consciousness. In N. J. Block, O. J. Flanagan, & G. Güzeldere (Eds.), *The nature of consciousness: Philosophical debates*. Cambridge, MA: MIT Press.

Rosenthal, D. M. (2003). Unity of consciousness and the self. *Proceedings of the Aristotelian Society, 103*(1), 325–352.

Schaffer, J. (2010). Monism: The priority of the whole. *Philosophical Review, 119*(1), 31–76.

Searle, J. (2005). Consciousness: What we still don't know. *New York Review of Books, 52*(1).

Shimojo, S., & Shams, L. (2001). Sensory modalities are not separate modalities: Plasticity and interactions. *Current Opinion in Neurobiology, 11*(4), 505–509.

Shoemaker, S. (2003). Consciousness and co-consciousness. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation.* Oxford: Oxford University Press.

Shoemaker, S. (2007). *Physical realization.* New York: Oxford University Press.

Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences, 353*(1373), 1295–1306.

Tye, M. (2003). *Consciousness and persons: Unity and identity.* Cambridge, MA: MIT Press.

Uzquiano, G. (2006). The price of universality. *Philosophical Studies, 129*(1), 137–169.

Van Gulick, R. (2004). Higher-order global states (HOGS): An alternative higher-order model of consciousness. In R. Gennaro (Ed.), *Higher-order theories of consciousness.* Amsterdam: John Benjamin.

Zeki, S. (2003). The disunity of consciousness. *Trends in Cognitive Sciences, 7*(5), 214–218.