

Tim Bayne

Co-consciousness

*Review of Barry Dainton's 'Stream of Consciousness'*¹

The unity of consciousness is surprisingly elusive for something so familiar; attempts to describe or analyse the character of this unity are sometimes accompanied by suggestive but frustratingly vague metaphors, sometimes by implausible and exaggerated claims which ascribe to consciousness near-magical and certainly chimerical properties (p. 195).

I: Overview

While much that is written about the unity of consciousness does, as Dainton says, traffic in vague metaphors and exaggerated claims, Dainton's book is a superb example of sober thinking and meticulous attention to detail. *Stream of Consciousness* can be roughly divided into three projects, projects that are bound together by co-consciousness. In the present context 'co-consciousness' refers to the relation that experiences have when they are experienced together.² For instance, when one experiences a pain in one's foot while listening to music, these two experiences will typically be co-conscious parts of a single multi-modal experience, feeling-a-pain-in-one's-foot-while-listening-to-music.

Dainton's first project (chapters 2–3) is a defence of the view that co-consciousness is a primitive, unanalysable relation. Dainton examines a number of the leading analyses of co-consciousness and finds them all wanting. Next, (chapters 5–7) Dainton defends the claim that co-consciousness accounts for the continuity of consciousness — its unity through time. A large part of this project involves developing and defending a model of the specious present. (Buried between these two projects, in chapter 4, is a tentative defence of the claim that synchronic co-consciousness is necessarily transitive: if e1 and e2 are

Correspondence: Tim Bayne, Department of Philosophy and Religious Studies, University of Canterbury, Private Bag 4800, Christchurch, New Zealand, Email: t.bayne@phil.canterbury.ac.nz

- [1] **Barry Dainton**, *Stream of Consciousness* (Routledge: London and New York, 2000), xvi + 254pp., ISBN 0-415-22382-2 (hbk)
- [2] Most psychologists (and some philosophers) use 'co-consciousness' to refer to the relation that two streams of consciousness have when they are co-instantiated in a single animal; this is a very different relation.

Journal of Consciousness Studies, **8**, No. 3, 2001, pp. 79–92

co-conscious, and e2 and e3 are co-conscious, then e1 and e3 must also be co-conscious.) Dainton's third project involves an examination of the ways in which experiences might be holistically related to each other. Dainton's discussion here is rigorous and illuminating — these chapters alone are worth the price of the book.

My comments will follow the general structure of *Stream of Consciousness* itself. However, I want to begin by saying something about Dainton's general approach. Dainton's method is exclusively phenomenological, in the broad sense of the term: he is only interested in the structure of consciousness as it is present in experience. This may sound paradoxical — how could consciousness have non-experiential structure? — but in fact many hold that experiences have non-phenomenal properties; physicalists, for example, think that experiences have physical properties. Dainton is not interested in the non-phenomenal structure of consciousness. When he refers to the spatial (and temporal) structure of consciousness he is referring to the structure of its content or character, structure that is phenomenally manifest.

II: The Unity of Consciousness: Synchronic Co-consciousness

What is co-consciousness?

What sort of a relation is co-consciousness? Is it a phenomenal relation? By this I don't mean to ask whether its relata are phenomenal items — clearly they are — but whether it is itself a phenomenal item. Dainton seems to think that it is: 'the fact that v1 [a visual experience] is co-conscious with a1 [an auditory experience] is clearly a phenomenal characteristic, something manifest in consciousness' (p. 215, emphasis in original). Again, 'There is no denying that our streams of consciousness do display a distinctive sort of unity, and this unity does not just consist in a relationship between certain experiences, it consists in a relationship between experiences that is itself experienced' (p. 4; cf. p. 88).

But can co-consciousness be intuited? This question prompts another: what would it be to intuit co-consciousness? Towards the end of his discussion of the unity of consciousness Christopher Hill wonders whether there is a form of co-consciousness that is pure, that 'has no distinguishing characteristics other than its ability to unite sensations':

Although at one point in my reflections on unity of consciousness I was strongly inclined to think that there must be a ghostly form of co-consciousness that answers to this description, I now feel that this view is wrong. It isn't possible to find this ghostly form of co-consciousness within one's experience (1991, p. 239).

Is Dainton's co-consciousness a spectral relation, a relation that has no distinguishing characteristics other than its ability to unite sensations? And if so, is there any such relation? I'm not clear on the answers to either of these questions, but perhaps the following will be of some help.

Consider the property of being someone's relation. One cannot merely be related to a person, rather, one must first be their sister, brother, aunt or the like.

‘Co-familiarity’ is a determinable of which particular familiar relations are determinates. Perhaps being co-conscious is akin to being co-familiar. Perhaps two experiences are only ever co-conscious in virtue of bearing a more determinate relationship to each other and there is no such thing as bare co-consciousness. Call this the ‘determinable model’ of co-consciousness.

Does Dainton endorse the determinable model? It’s not clear. At one point he suggests that spatial relatedness is a mode of co-consciousness: ‘phenomenal items are co-conscious when they are experienced together, but there are different ways for phenomenal items to be experienced together; being experienced as occurring in some spatial relation is one such way’ (p. 225). But there is no suggestion in *Stream of Consciousness* that Dainton endorses the determinable model as such. Should one endorse it? This is a difficult question. Anyone wishing to defend it must account for the fact that a wide variety of experiences — including experiences that lack spatial content — can be co-conscious. What relation could a mood and a thought have in common such that they are experienced together other than pure co-consciousness? One option would be to take synchronic and diachronic co-consciousness as temporal modes of co-consciousness. Perhaps ‘occurring at the same time’ and ‘occurring just afterwards’ are temporal determinables of co-consciousness in the same way that ‘being to the right of’ is a spatial determinable of co-consciousness.

Leaving aside the questions raised by the determinable model, let me turn to the stated goal of chapters 2 and 3, which is to establish that co-consciousness is a primitive, unanalysable relation. Dainton examines three accounts of co-consciousness and finding them all unsatisfactory concludes that ‘we seem to be left with only one alternative: we simply accept that diverse experiences can occur together, as co-conscious . . .’ (p. 84; cf. 216). The critical question here is whether these three accounts of co-consciousness exhaust the alternatives to Dainton’s ‘no-theory theory’. I’m not convinced that they do.

We can divide accounts of the unity of consciousness into two broad classes: ‘objectivist’ and ‘subjectivist’ accounts. Objectivists allow that co-consciousness can be at least partially accounted for in terms of factors that are not accessible from the first-person perspective, while subjectivists hold that co-consciousness can only be explained in terms of factors that are first-person accessible. Subjectivist theories can be further divided into two sub-classes depending on whether they advert to the content of consciousness in accounting for co-consciousness (content subjectivism) or whether they look beyond content in accounting for co-consciousness (non-content subjectivism).³ (Content subjectivism is not to be confused with the view that intentional content is ‘narrow’ — that it does not involve factors external to the agent. It is a claim about the analysis of co-consciousness rather than the analysis of content.)

Dainton examines three accounts of co-consciousness: introspection accounts, awareness accounts, and spatial accounts.⁴ All three are subjectivist accounts:

[3] I thank Barry Dainton for some helpful comments on this distinction.

[4] These are really families of accounts, for they can each be developed in a number of different ways.

they appeal to phenomenal factors in order explain co-consciousness. The spatial account is a version of content subjectivism. It appeals to the unity of the spatial content of experience to account for co-consciousness. (In effect, the spatial account holds that spatial co-consciousness is the only determinable of co-consciousness.) The awareness and introspection accounts of co-consciousness can be developed in either content or non-content terms, depending on how awareness and introspection are understood. If these notions are thought to involve non-contentful acts of attention then they qualify as non-content theories. If, on the other hand, awareness and introspection are construed as involving higher-order representation of experiences — as they often are⁵ — then these accounts qualify as versions of content subjectivism.

Although Dainton's discussion of subjectivist accounts of co-consciousness is thorough, he devotes little attention to objectivist accounts. This is somewhat surprising given that some of the more prominent discussions of co-consciousness in the recent literature are objectivist. Both Shoemaker (1996; 2001) and Hurley (1998) suggest that co-consciousness can be at least partially accounted for in functionalist terms, although they appeal to different sorts of functional roles. Dainton's failure to explore objectivist accounts of co-consciousness is not in and of itself objectionable, but it is problematic given his claim that we must take co-consciousness as a primitive because there are no viable accounts of it.

Even if Dainton is right to endorse a no-theory theory of the unity of consciousness, it is not clear that his no-theory theory is the best no-theory theory. Dainton's approach to the unity of consciousness seems to be 'bottom-up': he builds fully unified streams of consciousness out of particular experiences and relations of co-consciousness. Alternatively, one could take total experiences and the relation of subsumption as one's primitives. (Roughly, one experience subsumes another when the former entails the latter.)⁶ We might think of the subsumption model as a 'top-down' approach to co-consciousness. It is not obvious which of these models is more plausible; indeed, it is not at all clear that these two models are really all that different. Both posit a primitive phenomenal binding agent — co-consciousness in the case of the bottom-up view, subsumption in the case of the top-down approach — and they appear to be largely inter-translatable. Nevertheless, they may have different implications for such issues as the transitivity thesis (see below), and there may be reason to prefer one approach to the other.

The 'just-more-content' objection

Hurley (1998) has claimed that any subjectivist account of co-consciousness faces the *just-more-content* (JMC) objection. The objection is that subjectivism generates a (vicious) infinite regress: 'how can anything internal to content determine unity, given that content presupposes unity? What prevents the problem of

[5] See Hill's (1991, p. 231) comments on Parfit's awareness account of co-consciousness.

[6] David Chalmers and I discuss the possibility of one experience entailing another without subsuming it in Bayne and Chalmers (2001).

co-consciousness from applying all over again to it?’ (1998, pp. 70f.)⁷ Although Dainton doesn’t mention Hurley’s argument, he is aware of this line of thought. Indeed, he himself uses the JMC argument against certain models of co-consciousness.

The fact that v1 [a visual experience] is co-conscious with a1 [an auditory experience] is clearly a phenomenological characteristic, something manifest in consciousness. . . . Yet the co-consciousness of v1 with a1 is not a separate and additional experience, an experience over and above v1/a1. If we were to subscribe to the awareness-content model, there would be an additional component, a separate awareness of v1 occurring with a1, but we have already rejected this model of consciousness. If there were an occurrent phenomenal judgment that v1 is co-conscious with a1, there would be a third component in the total experience, but by hypothesis there is not (and if there were, the same issue would arise concerning the co-consciousness of the phenomenal judgment with the audio-visual experience) (p. 215).

Although Hurley couches the JMC objection in terms of content, it is not clear that non-content versions of subjectivism escape the objection. The heart of the JMC objection is this: how can anything internal to consciousness unify consciousness without itself having to be unified by a further phenomenal item? Once the objection is put in this very general form, it is no longer obvious that (or how) Dainton avoids the objection.

It is possible, I think, to read Dainton as making two responses to the JMC objection. First, he emphasizes that although co-consciousness is a phenomenal item, it is not itself an experience in its own right, it is, rather, an experiential *relation*. Second, Dainton holds that co-consciousness is a *self-binding* experiential relation. We don’t need to introduce R1 in order to bind R and e1 and e2 together, R not only binds e1 and e2 together, it also binds itself to e1 and e2 (p. 215). These two responses are independent, in the sense that (if successful) they are individually sufficient to block the JMC objection.

It is not easy to know what to make of the first response, for it is rather unclear what the difference between an experience and an experiential item comes down to. In order to generate the regress, all the proponent of the JMC argument needs is the claim that any phenomenal item will be co-conscious with every other phenomenal item that occurs within the same complex experience, and it is not clear that (or how) Dainton can avoid granting this point. Dainton’s second response to the JMC argument is more compelling. Of course, it supposes that there are self-binding phenomenal relations, but there’s no obvious reason (that I can see) for rejecting such a notion. Note that *content* subjectivists can adopt a

[7] At certain points Hurley gives the JMC argument a slightly different spin. She says that any attempt to unify conscious states in term of content faces the problem that any such content can be duplicated in separate centres of consciousness. My conscious states cannot be unified by my first-person thoughts, because someone else can have type identical first-person thoughts (1998, p. 61). I confess that I find it difficult to see what the problem is here. If experiences can unify they do so as tokens not as types. If first-person thoughts succeed in unifying experiences, then my first-person thoughts unify my experiences and your first-person thoughts unify your experiences.

similar response to the JMC argument by insisting that the content-bearing states that unify experiences also represent themselves as unifying experiences.

The combinatorial objection

Dainton's account of co-consciousness seems to be vulnerable to a further worry.⁸ Dainton holds that when experiences are co-conscious they fuse to form a state that is an experience in its own right (pp. 105, 215). So two experiences, e1 and e2, will fuse to form a third experience, e3. Since e1 and e2 are parts of e3, any total experience that includes e3 will also include e1 and e2. Furthermore, Dainton defines a total experience as a group of experiences that are mutually co-conscious (p. 95). It seems to follow from this that e1 and e3 will be co-conscious, as will e2 and e3, generating conscious states e4 and e5 respectively. But of course, e5 will be co-conscious with e1–e4, and so on, generating still further states of consciousness, which will themselves combine with their parts to form still further experiences. It appears as though there cannot be a total state of consciousness: take any state of consciousness with parts, one can construct a more complex experience by combining it with one or more of its parts. It seems to follow that any stream of consciousness that has more than one experience at a time will have an indefinite number of experiences at a time. This would seem to be a conclusion that Dainton would want to resist. Call this the 'combinatorial objection'.

There are a number of possible replies to the combinatorial objection. One response would be to hold that states that are formed by the co-consciousness of other states cannot themselves enter into co-consciousness with other states. This seems a rather unsatisfactory response. Apart from the fact that it would be purely *ad hoc*, it would entail that any stream of consciousness with three or more primitive experiences at a time would fail to be fully unified: two of its primitive experiences would combine to form a compound experience, but this compound experience would not be co-conscious with a third primitive experience. A second response would be to say that an experience is not co-conscious with any of its proper parts. But this objection is no more attractive than the first one. For one thing, Dainton holds that 'there is no one "right" way to divide a subject's overall experience over a given interval into parts. As is plain, no matter which division is considered, all the relevant parts are related by co-consciousness ...' (p. 84). This suggests that any experience is co-conscious with all of its parts, for presumably there is a division of the experience according to which both the experience and its parts count as experiences in their own right.

Dainton himself suggests that the simplest way to solve the problem is to deny that e4 and e5 (and similar constructs) are extra or additional experiences; the resultant state is just e3 under a different name.⁹ In other words, the combination of an experience with one of its parts just is the experience itself. This response has a certain phenomenological plausibility, but it does, I think, raise the question

[8] This line of thought was prompted by one of Lockwood's criticisms of Hurley's discussion (Lockwood, 1994).

[9] Personal communication.

of what sort of work can be expected of co-consciousness. Dainton seems to regard co-consciousness as a *constructive* relation, a relation that enables one to build complex experiences out of (numerically distinct) simpler experiences. This constructive role for co-consciousness seems to be somewhat threatened if certain numerically distinct experiences, namely those that stand in the part–whole relation, can be co-consciousness without the product of this relation resulting in an experience which is greater than either constituent of the relation.

The transitivity thesis

In his seminal paper on the effects of commissurotomy operations Nagel (1971) suggested that it might be possible for a person to have somewhere between one and two consciousnesses. Lockwood (1989) translated Nagel's suggestion into a proposal about the logical structure of co-consciousness. He suggested that some consciousnesses might be partially (or weakly) unified, where a partially unified consciousness is one in which co-consciousness fails to be transitive.¹⁰ The notion of partial unity is at odds with the intuitive claim that (synchronic) co-consciousness is transitive — a claim that we can call the 'transitivity thesis'. Should we jettison the transitivity thesis and conclude that partial unity is possible? Dainton thinks not. Although he argues that diachronic co-consciousness is not transitive — indeed, he suggests that diachronic co-consciousness is *essentially* non-transitive — Dainton argues that synchronic co-consciousness must be transitive. Although I am sympathetic to the transitivity thesis, I find Dainton's defence of it problematic.

The first problem concerns the difference between diachronic and synchronic co-consciousness. The worry here isn't just that Dainton holds that the former isn't transitive while the latter is; rather, it concerns the very distinction between the two. Dainton claims that the moment over which co-consciousness is necessarily transitive isn't an instant but is in fact 'some brief interval that is shorter than the specious present' (p. 172). In other words, Dainton holds that co-consciousness *is* diachronically transitive, albeit very briefly. But if transitivity in co-consciousness can hold across short durations why does it fail across longer durations?

A related problem undermines Dainton's argument for the transitivity thesis. His argument turns on the claim that when two experiences are co-conscious they are fused, they are wholly joined such that there is no 'distance' separating them. 'Since e1 and e2 are parts of a single experience in this way, how could it be possible for another experience e3 to be co-conscious with e2 without also being co-conscious with e1?' (p. 105). Those who reject the transitivity thesis will simply deny that fusion entails the transitivity of co-consciousness. Furthermore, on Dainton's own view transitivity can't be a *necessary* consequence of fusion, for Dainton accepts that when e1, e2 and e3 fall outside a single moment the fusion of e1 with e2 and e2 with e3 does not entail the fusion of e1 with e3.

[10] More precisely, a (merely) partially unified consciousness is one in which every experience is co-conscious with at least one other experience, but there are at least two experiences which fail to be co-conscious (see Lockwood, 1994, p. 91).

A third difficulty with Dainton's defence of the transitivity thesis is that it is open to what I call the 'duplication objection'. Suppose that we follow Dainton and individuate experiences in terms of their time, their exact phenomenal character, and their physical basis (pp. 24, 189). Such an account of experiences seems to permit us to make the *duplication assumption*, viz., that a single subject can have two experiences with qualitatively identical phenomenal characters at the same time.¹¹ Take such a subject, *S*, who at a certain time has three experiences, *e1*, *e1** and *e2*, where *e1* and *e1** are identical in phenomenal character. Suppose that *e1* and *e1** are co-conscious, *e1* and *e2* are co-conscious, but *e1** and *e2* are not. Although transitivity fails in this scenario, from *S*'s perspective it seems to hold. The crucial point is that introspection doesn't allow *S* to distinguish between *e1* and *e1**, and thus she can't tell whether *e1* or *e1** (or both) is co-conscious with *e2*. The transitivity thesis is defined over *token* experiences, but, as Hurley (1998) points out, there is a sense in which the subjective perspective only has access to relations of co-consciousness between experiential *types*.

There are two points at which the duplication objection is vulnerable to attack. First, one could question the duplication assumption. Perhaps the most straightforward way to do this is to individuate experiences in terms of their subject, so that a subject couldn't have two experiences with exactly the same phenomenal character at the same time. A second response to the duplication objection is to deny that duplicate experiences — that is, experiences with exactly the same phenomenal character that belong to the same subject at the same time — would, from the subject's perspective, coalesce into a single experience. I'm inclined to think that this second response cannot be sustained — and that Dainton cannot reject the duplication objection without rejecting the duplication assumption — but I lack the space to motivate this claim here.

III: The Continuity of Consciousness: Diachronic Co-consciousness

I turn now to the continuity of consciousness. Dainton has two projects here. First, he develops a model of the specious present according to which a specious present is a temporally extended total experience, that is, it is an experience that is composed of mutually co-conscious experiences with temporal content (p. 168). Second, Dainton develops an *overlap* model of the extended continuity of consciousness: successive specious presents belong to a continuous stream of experience because they overlap, i.e., they have (proper) parts in common. The stream of consciousness composed of *Do-Re-Mi* might involve one specious present (*Do-Re*) being replaced by a second specious present (*Re-Mi*), where these two specious presents have the (particular) experience *Re* in common. Dainton's analysis of phenomenal time is complex and sophisticated, and I can only touch on some of the many issues that it raises.

[11] The title of the objection and the thesis on which it is based is mine but the objection itself is Hurley's (see Hurley 1998, p. 105ff.) Don't confuse what I'm calling the duplication assumption with what Hurley (1998) calls the duplication assumption.

The specious present

Why posit a specious present in the first place? One motivation for the notion is theoretical: some argue that one cannot explain phenomenal temporality without it. Dainton endorses this defence of the notion, and has much to say about the problems confronting theories of phenomenal temporality that attempt to do away with a specious present. A second motivation for the specious present is phenomenological: some claim that one can be directly aware of temporal extension within experience itself, in much the same way that one can be directly aware of spatial extension within experience. Dainton also endorses this justification of the specious present. He estimates that the typical extent of his specious present is on the order of half a second or less (p. 171). Contrast this modest proposal with William James' remark that 'the specious present has . . . a vaguely vanishing backward and forward fringe, but its nucleus is probably the dozen seconds or less that have just elapsed' (quoted in Gallagher, 1998, p. 18). One might well wonder whether Dainton and James are talking about the same thing.

But what *is* the specious present? Dainton makes a number of claims about the specious present, some of which seem to hang together rather uncomfortably. On the one hand he says that a specious present is a total experience (p. 172), where a total experience is a fully unified experience (every part of it is mutually co-conscious with every other part of it) that is not a part of any larger fully unified experience (p. 96). At the same time, Dainton allows that the length of a single individual's specious present might be modality specific (p. 173). This prompts the following question: what happens to the content of a total experience when two modalities have specious presents of different lengths? The span of the total experience is only the length of the shorter specious present: 'Only then would every part of a total experience be co-conscious with every other part' (p. 173).

Given Dainton's definition of a total experience this is the appropriate thing to say, but it seems to lead to a strange state of affairs. Take a subject, *S*, who has an auditory experience (*a*1) and a visual experience (*v*1). Suppose that *S*'s auditory specious present is longer than her visual specious present. Take two earlier auditory and visual experiences, *a*2 and *v*2, that had been simultaneous, where *S*'s specious present includes *a*2 but fails to include *v*2. Although *a*2 is co-conscious with both *a*1 and *v*1 — it must be, because it is part of *S*'s auditory specious present — it isn't part of *S*'s current total experience, because it falls outside the scope of *S*'s shortest specious present. But since *a*2 is co-conscious with the rest of *S*'s current experiences — at least, it is co-conscious with *a*1 — why shouldn't it be included within (subsumed by) *S*'s current total experience? Surely *a*2, unlike *v*2, is part of what it is like to be *S* at this moment. This suggests to me that the length of the specious present cannot be modality-dependent.

Dainton argues that there is a third aspect to the temporality of consciousness over and above the phenomenal continuity that is internal to a specious present and the phenomenal continuity that extends across specious presents. He posits this third form of phenomenal temporality because he claims that temporal flow cannot be accounted for in terms of (diachronic) co-consciousness.

The fact that co-consciousness is symmetrical with respect to time is quite compatible with experience itself possessing an inherent direction: all that is required is for co-conscious experiences to have contents that are not symmetrical with respect to time. And clearly, the contents of our experience have this feature (p. 176).

So there are two forms of temporality internal to a single specious present. There is the fact that all experiences within a specious present are diachronically co-consciousness, and there is the fact that diachronic co-consciousness has directionality. In the same way that spatial co-consciousness can be asymmetrical, as when one experiences is enclosed within another, temporal co-consciousness can also be asymmetrical, as when one experience is prior to another. Indeed, there is some reason to think that temporal co-consciousness, unlike spatial co-consciousness, is necessarily asymmetrical: diachronic co-consciousness differs from synchronic co-consciousness precisely because the experiences so related are asymmetrically ordered in time.

Extended temporal structure

Dainton's analysis of extended temporal structure is intimately integrated with his account of the specious present, in that his overlap model of phenomenal continuity seems to entail that any temporally extended consciousness will have a specious present of some duration. If e_1 and e_2 are diachronically co-conscious, then they must generate a specious present of length $e_1 - e_2$. In what is a rather tidy result, the specious present emerges as a necessary consequence of phenomenal continuity.

Unfortunately, the overlap model is open to objection. In fact, it is vulnerable to a version of the duplication objection that I raised against Dainton's defence of the transitivity thesis. Consider a subject (S) who has two token Re experiences at a time, one (Re_1) is co-conscious with Do (and only with Do), the other one (Re_2) is co-conscious with Mi (and only with Mi). So, the subject's specious presents follow the following pattern: $\langle Do, Re_1 \rangle$, $\langle Re_2, Mi \rangle$. Are these two specious presents continuous? No, because they don't contain any experiences in common. Would they seem to S to be continuous? It seems to me that they may well do. If S only has introspective access to the phenomenal character of her experiences, then she won't be able to distinguish Re_1 from Re_2 on the basis of introspection. Dainton could block this objection if he held that any experience that is concurrent (and co-subjective) with experiences that belong to the same specious present itself belongs to that specious present (in which case S would have specious presents of $\langle Do, Re_1/Re_2 \rangle$, $\langle Re_1/Re_2, Mi \rangle$). As I've already suggested, perhaps Dainton should adopt this view.

But rather than dwell on Dainton's overlap model, I want to explore an issue that he touches on only in passing. Consider the following claims:

A typical stream of consciousness is a succession of experiences which lasts for some hours (p. 113; cf. p. 90).

[J]ust as a stream consists of an uninterrupted flow of water from start to finish, the stretches of our conscious lives *that span periods of dreamless sleep* consist of an uninterrupted flow of experience' (p. 117, my emphasis).

How are experiences that are separated by dreamless sleep, coma, or the like related? Is the continuity of a stream of consciousness necessarily interrupted by discontinuities in objective clock-time? Can a stream of consciousness survive an interruption in its phenomenal continuity? It is far from obvious that dreamless sleep does interrupt the continuity of consciousness. As Dainton points out (p. 131), there is no obvious reason why it should; after all, if the continuity of phenomenal space is not troubled by fissures in the spatial structure of the vehicles of consciousness, why couldn't the continuity of phenomenal time straddle discontinuity in objective time? It might be that Dainton's reluctance to assume that streams of consciousness span dreamless sleep has a phenomenological ground. Perhaps he holds that the first moment of consciousness after sleep is not phenomenally continuous with the last moment of consciousness prior to sleep. Note that on Dainton's model, the phenomenal continuity of this last experience before falling into a dreamless sleep (e1) and the first experience upon waking (e2) would imply that e1 and e2 form a single specious present (after all, they would be directly co-conscious). Can phenomenal continuity in general, and the specious present in particular, straddle dreamless sleep? These questions are among the many fascinating questions that Dainton's account raises in passing.

IV: Phenomenal Interdependence

In the final two chapters of *Stream of Consciousness* Dainton turns to the issues of phenomenal holism and interdependence. On the one hand Humeans claim that experiences are atomistic: like bricks in a wall, their character is unaffected by their phenomenal context. Holists, on the other hand, claim that the character of particular experiences is invariably affected by their phenomenal context. As William James put it, 'the most dim shade of perception enters into, and in some infinitesimal degree modifies, the whole existing state' (James, quoted in Dainton p. 184). Although Dainton's sympathies lie with Hume, he doesn't reject holism entirely.

Holism takes many forms; Dainton's focus is on what he calls the *complete phenomenal interdependence* thesis (CPI). According to CPI experiences take on the character of the entire stream of consciousness to which they belong, so that if any part of that stream had been different in any way that experience would have had a different intrinsic character. Since the character of an experience is essential to it, CPI entails that experiences necessarily occur in streams that are phenomenally indistinguishable from those in which they actually occur. As Dainton points out, CPI is a rather implausible position: 'Would a sense of melancholy alter in felt quality if the noise of the passing cars were a fraction louder, or the room slightly brighter? ... It seems unlikely' (p. 194). I concur whole-heartedly.

Dainton has more sympathy for the thesis of *partial phenomenal interdependence* (PPI). He grants that there is a limited amount of phenomenal

interdependence, although it is almost exclusively intra-modal: ‘there is little or no trace of inter-modal interdependence to be found at all’ (p. 195). On this point I part company with Dainton. Inter-modal interdependence may not be ubiquitous, but it certainly isn’t rare. The study of inter-modal effects is an established part of psychology (for a review see Stein and Meredith, 1993; for some particularly interesting effects see Bertelson, 1998; Botvinick and Cohen, 1998; Jordan and Bevan, 1997; Ramachandran and Hirstein, 1997). Consider just one such effect, the McGurk effect. In their seminal study, McGurk and MacDonald (1976) dubbed the sound of someone saying [ba] onto the lip movements for [ga]. Normal adults reported hearing [da]. Note that this is what subjects reported *hearing*. Visual information is involved in forming an *auditory* percept. ‘By merely closing the eyes, a previously heard [da] becomes [ba] only to revert to [da] when the eyes are open again’ (1976, p. 747). Ventriloquists rely on a similar inter-modal effect for their livelihood.

Most of Dainton’s discussion of phenomenal interdependence is focused on gestalt phenomena. He structures his discussion around the distinction between strong impingement and weak impingement (p. 199):

Strong Impingement: Phenomenal wholes have certain parts that possess intrinsic phenomenal features that reflect the character of the whole, and parts with the same character could not possibly occur except in a whole of the same or similar type.

Weak Impingement: The character of the constituent parts of a phenomenal whole are partly dependent on their being such, but items with just the same intrinsic phenomenal character as these parts could exist in wholes of a different type, or as perceived wholes in their own right.

Dainton argues that gestalt phenomena can be accounted for in terms of (contingent) Weak Impingement. Consider the Müller-Lyre illusion (*Figure 1* here; Figure 8.3 in Dainton’s book):




Figure 1. Müller-Lyre illusion

As Dainton points out, the orientation of the ‘fins’ makes a difference to how (most people) perceive the lines. But this only establishes Weak Impingement, for ‘it is not as though we cannot see lines of the same apparent phenomenal length in different contexts, e.g. where the ‘fins’ are absent’ (p. 199).

In cases of Weak Impingement, there is no enigmatic interpenetration of whole with part. Although how the part (regarded as an external stimulus) is perceived is affected by context, the effect of context on character is not so far-reaching that a phenomenally indistinguishable part could not be experienced in a different context (p. 201).

The crucial point here is the distinction Dainton makes between the external stimulus and the experience thereof. Take a cluster of stimuli, $s_1 \dots s_5$. One’s experience of s_3 (e_3) may depend on the nature of the stimuli surrounding s_3 . If s_3 were surrounded by qualitatively different stimuli, s_1^* , s_2^* , s_4^* and s_5^* , one might experience e_3^* instead of e_3 . We can describe cases such as these as involving phenomenal interdependence — as Dainton does — but they don’t quite live up to this label. ‘*Phenomenal interdependence*’ would be best applied to cases in which e_3 could only be perceived in the context of $e_1 \dots e_5$. This is exactly what happens in cases of Strong Impingement.

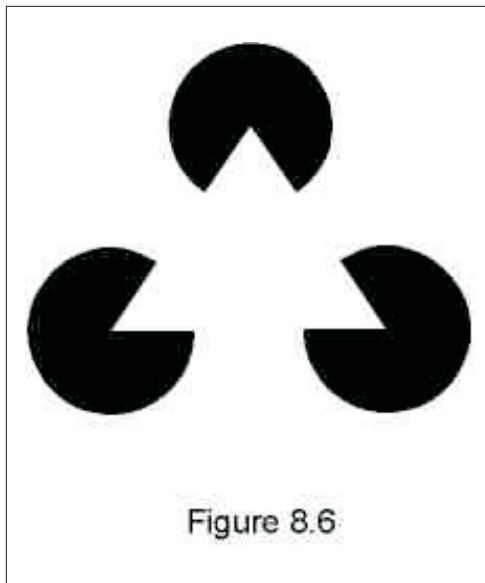


Figure 2. Kanisza's triangle

Although Dainton argues that there are few plausible examples of Strong Impingement, he does suggest that Kanisza’s triangle (*Figure 2* here; *Figure 8.6* in Dainton’s book) may qualify as such a case.

‘The triangle is an integral part of this phenomenal whole, and is clearly dependent upon the other parts: mask off two of the “pies” and the triangle vanishes; leaving all three “pies” in place but altering their orientation produces the same result’ (p. 206). True, but why does that show that Kanisza’s triangle is an instance of Strong Impingement? ‘What suggests Strong Impingement in this case is the difficulty of envisaging a phenomenally indistinguishable triangle occurring in a markedly different

phenomenal whole’ (p. 206). Is this right? Couldn’t this white triangle experience occur in the context of a wide variety of phenomenal wholes? Cut around ‘the white triangle’ that you see on this page and superimpose it on almost any solid background: aren’t you now having a type-identical white triangle experience in a markedly different phenomenal context? I suspect that Strong Impingement may even be rarer than Dainton thinks.

Conclusion

There is a wealth of material in *Stream of Consciousness* that I haven't had the space to mention: Dainton has interesting things to say about introspection, the act-object conception of experience, and phenomenal space. I hope, however, that what I have said suffices to generate the interest in this book that it deserves. Dainton says that his results 'have not been negligible or without interest' (p. 236); I would suggest that this is something of an understatement.¹²

References

- Bayne, T. and Chalmers, D.J. (2001), 'What is the unity of consciousness?', in *The Unity of Consciousness: Binding, Integration, Dissociation*, ed. A. Cleeremans (Oxford: Oxford University Press).
- Bertelson, P. (1998), 'Starting from the ventriloquist: the perception of multi-modal events', in *Advances in Psychological Science, vol. 2: Biological and Cognitive Aspects*, ed. F. Craik and M. Sabourin (Hove: Psychology Press/Erlbaum).
- Botvinick, M. and Cohen, J. (1998), 'Rubber hands "feel" touch that eyes see', *Nature*, **391**, p. 756.
- Gallagher, S. (1998), *The Inordinance of Time* (Evanston, IL: Northwestern University Press).
- Hill, C. (1991), *Sensations: A Defense of Type Materialism* (Cambridge: Cambridge University Press).
- Hurley, S.L. (1998), *Consciousness in Action* (Cambridge, MA: Harvard University Press).
- Jordan, T.R. and Bevan, K. (1997), 'Seeing and hearing rotated faces: influences of facial orientation on visual and audiovisual speech recognition', *Journal of Experimental Psychology: Human Perception and Performance*, **23** (2), pp. 388–403.
- McGurk, H. and MacDonald, J. (1976), 'Hearing lips and seeing voices', *Nature*, **264**, pp. 746–8.
- Lockwood, M. (1994) 'Issues of unity and objectivity', in *Objectivity, Simulation and the Unity of Consciousness*, ed. C. Peacocke (Oxford: Oxford University Press).
- Lockwood, M. (1989), *Mind, Brain and the Quantum* (Oxford: Blackwell).
- Nagel, T. (1971), 'Brain bisection and the unity of consciousness', *Synthese*, **22**, pp. 396–413.
- Ramachandran, V.S. and Hirstein (1997), 'Three laws of qualia: what neurology tells us about the biological functions of consciousness', *Journal of Consciousness Studies*, **4** (5–6), pp. 429–57.
- Shoemaker, S. (1996), 'Unity of consciousness and consciousness of unity', in *The First-Person Perspective and Other Essays* (Cambridge: Cambridge University Press).
- Shoemaker, S. (2001), 'Consciousness and co-consciousness', in *The Unity of Consciousness: Binding, Integration, Dissociation*, ed. A. Cleeremans (Oxford: Oxford University Press).
- Stein, B.E. and Meredith, A.M. (1993), *The Merging of the Senses* (Cambridge, MA: MIT Press).

[12] I am very grateful to David Chalmers and Barry Dainton for their comments on earlier versions of this paper.