

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/228733044>

Saving the Baby: Dennett on Autobiography, Agency, and the Self

Article in *Philosophical Psychology* · July 2006

DOI: 10.1080/09515080600690565

CITATIONS

15

READS

89

1 author:



Jenann Ismael

Columbia University

70 PUBLICATIONS 759 CITATIONS

SEE PROFILE

Saving the Baby: Dennett on Autobiography, Agency, and the Self

Jenann Ismael

Dennett argues that the decentralized view of human cognitive organization finding increasing support in parts of cognitive science undermines talk of an inner self. On his view, the causal underpinnings of behavior are distributed across a collection of autonomous subsystems operating without any centralized supervision. Selves are fictions contrived to simplify description and facilitate prediction of behavior with no real correlate inside the mind. Dennett often uses an analogy with termite colonies whose behavior looks organized and purposeful to the external eye, but which is actually the emergent product of uncoordinated activity of separate components marching to the beat of their individual drums. I examine the cognitive organization of a system steering by an internal model of self and environment, and argue that it provides a model that lies between the image of mind as termite colony and a naïve Cartesianism that views the self as inner substance.

Keywords: Dennett; Autobiography; Self; Self-Models; Navigation; Self-Organization; Unity Of Self

1. Introduction

As I stand here, I keep a kind of running tab on myself and my situation in the world. I ascribe to myself various thoughts, feelings, hopes. I recall past events and think about my future. This first-personal stream of consciousness running through our lives—this continuing jumble of thoughts, experiences, and emotions, all self-ascribed and woven into a Joycean monologue in which *I* am both narrator and central character—is something we all have and take one another to have. Maybe no other animals have it; maybe we're the only ones that represent ourselves, and represent ourselves *as* experiencing subjects. Maybe we're the only ones, that is to say, that engage in reflexive self-representation. What is the point of it, what role does this

Correspondence to: Jenann Ismael, Centre for Time, Department of Philosophy, Main Quad A14, University of Sydney, NSW 2006, Australia. Email: jenann.ismael@arts.usyd.edu.au

self-representing inner life play in the intrinsic dynamics of the body? Probably nobody in the philosophical literature has talked about this more than Dennett, and his views are characteristically iconoclastic. He thinks this running narration is empty chatter: the “I” that is supposed to be the proper subject of experience, thought, and action doesn’t exist. In his view, autobiographies are a confabulatory byproduct of the decentralized brain activity that actually regulates behavior: someone who mistakes it for an accurate portrait of the source of behavior makes the same mistake as someone who mistakes the smoke spewed out by an engine as the power that drives it.

There is a genuinely new, and interesting form of anti-realism here—one that forces a confrontation between the decentralized view of human cognitive organization finding increasing support in parts of cognitive science and the centralization suggested by talk of an inner self. *You*, after all, are supposed to be the proper subject of thought, experience, and action; and if it turns out that there is no internal locus of knowledge and control, that will look like a reason to deny that there is a self. I will be embracing much of what Dennett says, but rejecting the conclusion.

2. The Cartesian Theater

The Cartesian interpretation takes the Joycean monologue at face value: the “I” of the inner monologue is an internal homunculus—the individual consciousness—that receives the information coming in through the senses and controls the activity of the body. Dennett thinks that cognitive neuroscience eliminates the self *qua* bearer of knowledge and controlling agent by revealing the complex, decentralized activity in the brain that is actually responsible for behavior. Close examination is supposed to reveal that there is no “brain pearl,” “Oval Office,” or “Cartesian Theater” in the brain. There is nothing but multiple processing streams, non-intersecting causal and informational pathways, a motley “bag of tricks” that neither requires nor supports the centralized vision of the autobiography. The “I” that is supposed to be choreographing the complicated ballet of bodily motion is a fiction. His suggestion is that if we look closely at the brain, the impression of centralized control, which is so irresistible from an external perspective, dissolves. He tends to put the case by analogy with what are known in the scientific literature as self-organizing systems—of which the termite colony is a central example—systems in which apparently coordinated activity arises from the joint operation of autonomous subcomponents. Just as evolution eliminated god from the natural world by providing a self-organizing explanation of biological design, the suggestion is that neuroscience eliminates the self by providing an account of how separate subsystems in the brain generate coordinated behavior without any central supervision.¹

We do know that nature is full of systems in which the impression of centralized control is misleading. These kinds of system have become the focus of a flurry of excited interest, and the practice has taken hold in fields that range from physics through the social sciences of treating any complex dynamical system

(i.e., any system that has parts and whose state changes with time) as a self-organizing one. Dennett's contention is that science reveals that the brain-body system is self-organizing, that the impression of centralized control is misleading, and that all of the talk about the inner "I" directing the activity is fiction. We produce autobiographies, in his view, for the same reason that we give centralized explanations of the behavior of termite colonies, traffic systems, and schools of fish though we know there is no such thing, namely, because it's a useful fiction that faithfully tracks the gross movements of the system and that abstracts from microfacts that introduce pointless, irrelevant mathematical complexity. In his words (Dennett & Humphrey, 1998):

The revisionist case is that there really is no proper-self: none of the fictive selves—including one's own firsthand version—corresponds to anything that actually exists in one's head . . . complex systems can in fact function in what seems to be a thoroughly "purposeful and integrated" way simply by having lots of subsystems doing their own thing without any central supervision. Indeed most systems on earth that appear to have central controllers (and are usefully described as having them) do not. The behavior of a termite colony provides a wonderful example of it. The colony as a whole builds elaborate mounds, gets to know its territory, organizes foraging expeditions, sends out raiding parties against other colonies, and so on . . . Yet, in fact, all this group wisdom results from nothing other than myriads of individual termites, specialized as different castes, going about their individual business—influenced by each other, but quite uninfluenced by any master-plan. (pp. 39–40)

Where Descartes looked for an inner homunculus, what science finds is a bit of machinery with the job of concocting a first-personal narrative using various contributions from different parts of the mind. The narrative is there to encourage the useful fiction of a central controller. We tell it to others, and get wrapped up in it ourselves, forgetting its fictional origin, and go on spinning it in blissful unconsciousness of the complex, decentralized activity in the brain that actually regulates behavior. The stream of consciousness is just a subvocalization of this fiction, a confabulatory product of the "Joycean Machine" in the head, with no role in the intrinsic dynamics of the body. The inner voice has no owner. The whole thing is a put up job.

3. Getting Rid of the Bathwater

It's easy to get lost in the welter of Dennettian metaphors. There are two that matter for our purposes: Joycean Machine and the Cartesian Theater. The Joycean Machine, as I understand it, is the machinery in the head that composes the autobiography and controls verbal reports in the first person. It's important to Dennett that this machinery is thoroughly distributed, involving multiple subsystems in the brain in an attenuated process of writing and rewriting. The Cartesian Theater is the fictional inner stage across which the events in our psychological self-portraits are supposed to march if, according to Dennett, we take them at face value. To reject the Cartesian

Theater in favor of the Joycean Machine is to see every part of one's brain as actively involved in composing an autobiography without any presumption that the progression of thoughts and impressions self-ascribed in the psychological part of that portrait actually play out in any separate arena.

There is much that we can, and should, agree with here. First, there is no question, as I said, that there are systems which look as though they are controlled by a central intelligence but in which the collective behavior is emergent from the unchoreographed dynamics of autonomous components. And certainly a lot of what we do, and a lot of what other animals do, can be done without any central controlling agent. Second, we should ruthlessly excise the Cartesian bathwater in our conceptions of our own minds, replacing the Cartesian Theater with the Joycean Machine, focus on how brains "grow self-representations," and reject the idea that there is anywhere other than the inner monologue itself where the activity it reports plays out. The brain is a storehouse of information brought in through different sensory pathways. Information that has a fixed bearing on behavior is used by sensorimotor subsystems without conscious registration. The Joycean Machine selectively culls information from this cacophony of non-conscious, non-intersecting causal and informational pathways and, in a cooperative and attenuated process, pieces together a coherent narrative, a self-centered portrait of an objective world as seen, felt, smelled, and heard from the perspective of the narrating subject ("here I stand, feeling thus and so, looking out at this and that, thinking of such and such . . ."). We should recognize, moreover, that the process of piecing together a coherent story is highly constructive, and the narrative it produces is gappy and prone to retrospective revision, getting filled in on the fly. Dennett's concern about the veridicality of this portrait is two fold. He thinks that it requires the existence of an inner locus of both information and control, and he thinks that there is neither.²

I believe there is an important and deeply anti-Cartesian insight in replacing the Cartesian Theater with the Joycean Machine and that if we really want to understand what selves are, we should be looking, as Dennett says, at how the brain produces self-portraits. I concur that those portraits are not pictures either of an object or arena that exists inside the brain, but products of a synthesis wrought by the brain as a whole. What I want to take issue with is Dennett's interpretation of the monologue as requiring the existence of an inner object, and his view about the role of monologue in the intrinsic dynamics of the body, if we take seriously the dynamical analogy with self-organizing systems and the repeated insistence that the monologue is a fiction fabricated for an external audience. What I'm going to suggest is that the autobiographical monologue pieced together by the Joycean Machine—the "Joycean stream," as I'll call it—has an important and substantial role in the intrinsic dynamics of the body, and that there is no reason to suppose that the content of the monologue (aside from the usual mistakes of fact) is anything but a perfectly veridical portrait of the world centered on the system that produces it. The psychological component of that portrait—the part that depicts the progression of thoughts, ideas, and impressions running through the mind—should be interpreted reflexively, as a description of the stream itself.³

All of this is easier to see with an example, and I'm going to quite deliberately consider a system that Dennett himself suggests, not unfacetiously, as a model of self-representation: a ship that steers by a map.⁴ By 'ship' I don't mean just the physical vessel, I mean the whole complex system including the crew and the instruments and the computer networks that support it. I will suppose that navigation goes in cycles; sightings are made and instrument readings are taken, these are transformed by an attenuated chain that involves representation and re-representation across various media into a fix on the ship's coordinates, plotted on an objective map of the surrounding waters. On modern vessels, the ship ordinarily leaves shore with a fully prepared map, so that its only epistemic task is to update its own position, but we can easily grant vessels a role—like early explorers—in composing the maps that they steer by. Whatever the division of labor, some combination of design, learning, and experience provides a ship with enough built-up structure that, by the time it's navigating, it has some rudimentary map,⁵ can transform instrument readings into a fix on its location on the map, and add to the objective content of the map as it goes. For convenience, we'll suppose that the map is stored in a computer, whose contents are displayed on a screen in the main cabin. This just makes it easy for us to monitor its contents from the outside. Once the ship's coordinates are plotted on the map—which we can suppose is done by placing a little model of the ship on the map, something with enough structure to orient it and that encodes perhaps other properties of the ship important to navigation—an officer looking at the screen, or a team of officers, or the computer itself, swings into action, comparing the plotted location against a targeted destination, using information embodied in the fixed content of the map to plot and compare alternate routes, and ultimately arriving at a decision of the form "head due north at 98 knots."⁶ The decision is fed back to the distributed crew who transform it into operational procedures designed so that, if the environment cooperates, the ship moves in the prescribed direction. The cycle then begins again. New instrument readings are taken, transformed, and fed to the computer. New coordinates are plotted and compared against a targeted destination. The map runs through the decision process, a course is chosen, fed to the crew, and translated into motion.⁷ We can suppose that crew members are automatons, transforming commands that they recognize syntactically into operational procedures with no understanding of the contents of the commands, no apprehension of the position of the ship in space, so that the only place that that information is represented is in the main cabin, and—if we eliminate the officers in the cabin and (putting aside reservations about computability) let the computer run through the decision cycle—on the map itself.

If we focus just on the evolving contents of the map, ignoring all of the activity that's not explicitly represented there, what we see is an informational stream whose content is that of an evolving, objective representation of the spatial landscape centered on the ship.⁸ This informational stream receives input from the environment in the form of informational states that have the contents of self-locating beliefs, but it is propelled by an internal logic that transforms those

states into prescriptions for action, and that internal logic has roughly the form of deliberation. Those prescriptions for action, moreover, feed back into the ship and, provided all goes well with the rest of the machinery, guide the movements of the ship.

We, too, steer by self-centered maps of a world cast in impersonal terms. Our maps are more complex and richer in content than those of most ships. They have a temporal dimension, a rich array of representationally significant qualitative properties, and causal structure. They're at least partly linguistically coded and they have an explicitly psychological component. We don't just monitor our spatial locations, we keep track of our physical properties and our representational states, described in explicitly intentional terms, and we incorporate all of it into our self-models, together with an explicit record of our personal histories. All of these elements can be reproduced on the ship by adding either to the objective content of the map, or the self-model that orients the system in the space represented by the map. And we can begin to understand in a schematic way why some bits of information should be explicitly represented, but not others, and why some behaviors might be usefully brought under deliberative control and not others. The role that this informational stream is playing is something like that of the CEO of a vast, and largely self-regulating bureaucracy: unaware of the day to day activities that keep the system running, but setting long term goals, keeping track of the system's progress, and exerting influence needed to nudge behavior in the direction of goals.

The Joycean Machine, as I understand it, is not the informational stream itself, but the unrepresented apparatus that supports it. It does the work that is done on the ship by the crew and supporting equipment that transforms instrument readings and so on into the self-centered representation displayed in the main cabin. The central display here, however, is entirely eliminable. It translates the informational stream into a form that is easily visible from the outside, to the likes of us, but for the ship's purposes, it could be buried in the complex inner workings of the computer, discernible only to subsystems who need to use the information it contains, and conveyed to them in a not necessarily visual form. The stream has to interface with the subsystems on the ship that it controls, and it can speak to them in their own language. In the dynamical economy of the ship, it's the map itself that acts as both locus of information and control, bringing together information drawn from different sensory subsystems, rendering it in a form in which it can be combined with information acquired at different times and places, and brought to bear jointly on behavior. It does a good deal more than paint a pretty picture for an external audience, it creates the representational space within which deliberation takes place, and (although Dennett is surely right that we're a lot less deliberative than we think), at least in some cases, the deliberative processes that it represents as guiding really do guide. The dynamical organization of a system whose behavior is guided by this sort of unified representational space is quite different from that of a termite colony, or even a Brooksian robot, in which behavior is the emergent product of the joint

operation of a collection of independent subsystems, each with its own representational space doing its own thing. To make the comparison explicit:

System 1: Self-organizing system.

Informational streams leading from disparate peripheral sources feed directly into separate motor systems. The apparently coordinated behavior of the system is emergent from the collective activity of autonomous sensorimotor components.

System 2: Navigating system.

Information from disparate peripheral sources is plotted on a central map and used to generate a self-centered stream in which it is passed through a deliberative cycle before being fed into the motor pathways that govern the movements of the body.

System 3: Dennett's model.

Information from disparate peripheral sources is plotted on a central map and used to generate a self-centered stream in which it is passed through a deliberative cycle, but the stream now empties into the environment. Behavior is still controlled by subsystems that bypass the deliberative processes.

Dennett's model occupies an uncomfortable, intermediate position. It recognizes the Joycean Machinery that unifies the informational streams leading from the sensory into the motor pathways, and even goes through the stages of deliberation, giving reasons for behavior, generating prescriptions for action, and representing those prescriptions as causally implicated in the production of motion. But it makes the stream an idle wheel in the internal dynamics. We can do away with the Cartesian Theater, embrace the constructive account of how brains "grow self-representations" and still salvage an inner locus of information and control by regarding ourselves as navigating systems. The informational hub, the place in the mind "where it all comes together," is not the mythical Theater, but the Joycean monologue itself, which acts also as the seat of agency. It's an easily recognizable kind of dynamical organization, contrasting both with the pure self-organizing alternative and Dennett's variant of it.

A Joycean Machine is an expensive piece of equipment, and Mother Nature doesn't ordinarily spend on equipment that doesn't earn its keep in dynamical benefits. What might have led Dennett, who has himself emphasized Nature's frugality in myriad ways, to (at least apparently) deny that the stream it issues plays the directorial role it self-ascribes in the production of behavior? Presumably, the example of self-organizing systems, the insights of embodied cognition, and the success of cybernetics in deploying those insights in designing systems that negotiate complex environments successfully without explicit map-keeping. These systems cleverly use direct informational links between sensory and motor pathways and feedback loops to generate emergent patterns of behavior that mimic some of the characteristics of navigating systems. They have limits, however, that emerge when we look not at their first-order responses to stimuli, but at how those responses to stimuli change with changes in situation. The dynamical advantages of self-modelling can be put in a nutshell, and are quite nicely illustrated with the example of the ship. Explicit representation of the relationship between self and environment

adds degrees of freedom that allow a system to adjust its responses spontaneously to changes in its situation independently of the occurrent stimulus. Map-keeping goes together with self-representation, and map-keeping allows a system to bring the full body of explicitly stored information to bear on situated action in a manner that depends immediately and directly on the values of self-locating parameters (for development of this point, see Ismael, in press).

In the end, I want to downplay the importance of whether or not we say that selves exist. I concur with Dennett that there are no brain pearls, and there is no separate Theater in the mind in which our psychological histories play out. I'm inclined to think that these are dispensable myths about the self, encouraged by the grammar of 'I', but the issue seems to me to hold little interest. All that there is, and all that is needed to support "I"-use, is the Joycean stream itself. What does seem to me of central importance is that we understand with some clarity the difference between a system whose movements are orchestrated in part by a unified, self-centered, informational stream, and a system like an ant colony or a Brooksian robot in which behavior is the emergent product of the joint operation of a collection of non-intersecting informational streams.

We need to steer a careful path between the inner theater and the ant colony. It is important that we can view the Joycean Machine as a late addition to cognitive architecture that is self-organizing at the foundation. If we want to view ourselves as naturally evolved creatures, we need to see a line of development that leads from simpler systems to self-modelers. And we don't want to *unlearn* what we've learned from designing robots and studying systems like termite colonies about how much can be done *without* explicit representation. What practical influence the Joycean stream exercises is commandeered by diverting causal pathways from other subsystems, and only behavior that can't be effectively controlled by subpersonal mechanisms is usefully diverted. But we shouldn't lose sight of the virtues of steering by explicit self-centered models. We can find a place for navigation without falling into a mistaken Cartesian view of the mind.

4. The Many Voices of Dennett

I interpret Dennett as denying that there is an inner locus of information and control. In contrast, I have argued that, in a system that is steering by a self-centered map, the map itself serves as locus of information and control, providing a representational space where the information that is coming in through sensory pathways is plotted together and brought to bear jointly on behavior. If Dennett sometimes speaks as though the deliberative reasoning that occurs in the self-centered space of the map has no role in the intrinsic dynamics of the body, he doesn't always speak that way. It's an artifact of the richness of his work that one can find multiple lines of argument, and when they can be pulled apart, it's not always easy to reconcile them. I have followed one of these here: the one that advocates the analogy with self-organizing systems, insists that information and

control are both thoroughly distributed and seems to deny that the Joycean Machine is anything but an idle wheel in the dynamical economy of the body. I don't want to play the role of Joycean Machine to the many voices of Dennett, but if I were trying to weave a consistent narrative, this would strike me as a rather discordant thread. When he is speaking in other voices, he assigns the Joycean Machine a central, and quite substantial, role. It becomes the repository for information, the global workspace in which information is deposited for system-wide use, the place where opportunistic coalitions are formed by separable subsystems.⁹ He even refers to the Joycean Machine as the "control center" of the human organism.

There are passages in *Consciousness Explained* (1992) and in *Freedom Evolves* (2003) that seem to be in agreement with what I have said about the dynamical role of the Joycean Machine. And Dennett has described better than anyone the bootstrapping process of self-regulation and internal control that leads to the person-making qualities that distinguish us from other creatures. But it's only within an explicit representation of self against the background of an objectively rendered world that the ability to take oneself and one's relation to the world as intentional objects, and the associated capacities for self-reflection and self-evaluation, arises. The stream of consciousness is the space within which self-reflection and self-evaluation are possible: the Joycean Machine creates that space.¹⁰

I don't know how to reconcile this tension. It's closely related to one that Clark (2002) has diagnosed between the tendency to think that human minds are entirely continuous with the minds of other adaptive creatures, and on the other to think that there is something special about the human mind, something that underwrites the matrix of quite specialized capacities that accrue to specifically *human* intelligence. Dennett identifies language as the source of these capacities, arguing that language *gives* us labels for states of ourselves and of the world, and allows us to take the relationship between them as explicit object of thought. But there are strong indications of a deeper architectural difference that makes language, except in the most rudimentary way, specifically available to the human mind. It's not just that other animals don't develop linguistic frameworks of the same depth and scope as our own, it's that we cannot export our language across species lines. We cannot teach even our closest phylogenetic relatives to use languages that we have developed.

The tension is resolved if we can see language as rooted in the development of explicit self-representation—not simply representation of the movements of our bodies through a spatially extended landscape, but representation of ourselves and our states in a causally structured world. This much richer representational environment is the context within which the conception of ourselves as perceivers and agents, effecting and effected by events in our surroundings arises. It is the one within which the distinction between our own states and states of the environment, and between our own actions and events in our surroundings, can be made out. And it is the one that is required for the bootstrapping process self-regulation and self-control that Dennett identifies as the source of the matrix of

person-making qualities. This accords the Joycean monologue a position of quite extreme importance: information hub, seat of reasoning, locus of control. Full-blown personhood arises, by Dennett's own account, out of the processes of self-evaluation and self-definition that occur within it.

5. Autobiographical Memory

Why do we represent our pasts? Not simply why do we represent *the* past, but why do we each represent our *own* pasts? Why don't we discard information about personal history as soon as it has been used to ascertain our current situation? If the story I just told explains the dynamical benefits of routing sensory information through a self-representational loop that generates an explicit representation of current situation in a spatiotemporally extended landscape, it doesn't explain why our *self*-representations have a temporal dimension. Go back to the map-keeping system on the ship. We could easily add a temporal dimension to its map, but still keep it tracking only its current location so that the reflexive, or first-personal, content of its representation remains "here I am now." We, by contrast, keep a dynamic record of personal past, written and rewritten with each passing moment. To say that we look at the past through the lenses of the present doesn't begin to capture the complexity; we look at it through the layered lenses of all of the presents that separate us from it, and each of these lenses has its own transforming effect. If we have learned anything from studying situated cognition it is that Mother Nature doesn't reify structure without purpose. The reification of structure along the temporal dimension is so extravagant, it cries out for explanation. It is important to understand just how extravagant that structure is; every momentary cross section of our internal lives contains a selective, backward-looking image of the whole, one that is not only updated, but subjected to ongoing, retrospective re-vision.

The explanation has the same form for maps of all kinds.¹¹ The point of the reification is to bring information into the deliberative loop where it can bear (collectively, and in a manner that is regulated by self-location) on behavior. So the question is: what role does information about our personal pasts play in deliberation? Stated in this form, the question is easy. A system that retains an explicit record of its past can make promises, accept and acquire commitments. It can form developing relationships and form personal projects. It has interests of its own and can systematically carry out parts of an extended plan.¹² All of these capacities depend on retaining information about personal past in an explicit form. Just as the actions of a ship steering by a map are guided not just by its immediate spatial environment, but by its relationship to targets that lie outside its current sensory horizon so the actions of my present self are guided not just by my immediate past, but by my more distant past, and by the contribution of these actions to a long-term goal.

6. Unity

Dennett expresses his view as a challenge to the reality of the self, but the underlying issue has to do with unity. His claim was, in effect, that the unity implicit in the notion of the thinking-perceiving subject dissolves on close inspection as the mind separates into a collection of autonomous subsystems. The example of complex systems that steer by maps reveals new types of unity, types of unity that distinguish them from primitive loci of mental life, on the one hand, and from termite colonies, and schools of fish, on the other. I want to say briefly what they are. A Joycean Machine plays an important unifying role. Its job is to recombine informational streams that were separated by passage through the body by mapping them into an internal reconstruction of the external environment. That internal reconstruction provides the representational context within which deliberation takes place. There are three related types of unity that emerge from this arrangement. The first is what I will call the “synthetic unity” attained when information drawn from incommensurate sources are mapped into a common frame of reference. The second is the “univocity” attained when a set of separate, potentially conflicting informational streams is united into a single, collective voice. The third is the “dynamical unity” achieved when the parts of a system operate under the command of this single voice.

7. Synthetic Unity

One doesn't combine informational streams by just having them dump into a common bin. The information has to be integrated, which is to say, it has to be sorted, and organized and plotted in a common frame of spatiotemporal and descriptive reference that identifies overlap, and relates information coming in through different pathways. This is the job of the Joycean Machine. It is not different in kind from the integration of information effected by a detective weaving a coherent narrative out of the partially overlapping testimony of witnesses. Personal time lines and spatial movements are plotted jointly in an objective frame of reference to identify points of contact. The time at which one witness was watching the late show is identified with the moment another saw a stranger in the alley, a bang heard by one is connected to a flash seen by another, a glove reported missing by one witness is identified with one found later at the scene. The story goes through multiple drafts, information is added where narrative holes emerge, trajectories are continued when they pass out of sight, choices are made to resolve conflicting accounts. The integrated story doesn't report something seen by any one witness. That is not its role: it is a reconstructed compilation that identifies points of contact and reveals the relations between them.

8. Univocity

This process of integration resolves the cacophony of voices in the brain into a single coherent stream. When you resolve a collection of potentially competing

informational streams into a single coherent stream, I will say you give them a “collective voice.” It is this resolution of voices that allows the system to act as a unified agent. Consider the way that elections and referenda turn the cacophony of competing individual voices in a population into a single coherent stream. A population is officially polled for answers to questions on matters of public importance and the results of polls act in an official capacity as the collective voice of the people. The electoral process *gives* the community a collective voice. We can speak in a loose and metaphorical way of the voice of the people, but without the electoral machinery that resolves the collection of individual opinions into a collective “yes” or “no,” there is no truly univocal opinion. Or consider the procedures that turn the opinions of Supreme Court justices into a single opinion, expressed in the form of written or spoken pronouncement issued by the collective. Without collective pronouncements, there is the collection of opinions, but no opinion of the collective. When we say “The people have chosen so and so as their new President” or “The court has decided that the amendment to the law is unconstitutional,” the attitudes in question are attributed to the collective. The choices of the people are not mine or yours, they are ours. The opinions of the court do not belong to Sutter or Roberts, they are property of the group.¹³ Collective opinion need not be any simple function of the opinions of its members. Indeed, in the typical case the process that generates the collective opinion—the back and forth of debate leading up to an election or a ruling—has a complex dynamics involving feedback and feedforward connections that changes the opinions of the group and makes it effectively unpredictable. When it is given a collective voice, a complex system constitutes an intentional system in its own right.

The Joycean Machine gives the mind a collective voice. Without it, there is a collection of informational streams leading from the sensory surfaces into the motor pathways, but no whole to which representational states are properly ascribed in anything but a derivative, metaphorical sense. With it, the mind is a unified representational system with a voice of its own. The attitudes self-ascribed in the Joycean monologue belong to the unified system as surely as those self-ascribed by the spokesperson for the Supreme Court belong to the justices as a group. Does a voice-piece report? I would rather say it asserts, with an emphasis on the performative character of assertion. Reportage carries the implication of a subject matter that obtains independently of its pronouncement. A voice-piece *makes* true what it reports *by* pronouncing it. The people haven’t spoken until election results are in, the Supreme Court hasn’t made a decision until that decision is self-attributed on behalf of the collective by its spokesperson, i.e., until the spokesperson says something like “we, the justices of the court, register the opinion to follow . . .” Performances of this sort are truth bearing but self-fulfilling.

Whereas Dennett treats the Joycean Machine as a propaganda artist, like a White House crony that misrepresents leaderless government as a unified agent, I treat it as a voice-piece. The “I” of the mind isn’t a brain pearl lodged in the wetware, any more than “we, the people” is a controlling agent lurking among the populace. Looking in the population for an owner of the collective voice, or searching on the bench for one

who makes the court's decisions, one will find oneself as bereft as Dennett looking in the brain for a self (or, indeed, Hume searching his impressions for it).

An important lesson of all of this is that there doesn't have to be a controlling intelligence *in* the system to support the attribution of intentional attitudes to it. If subjects exist, and they are not put into the world by hand at the fundamental level, they are going to have to be composed somehow out of more basic—and ultimately non-intentional—elements. The kind of unity that is possessed by subjects of intentional attitudes is not given, but achieved, and the suggestion here is that it is achieved by forging a collective voice. Integration of the informational streams leading from sensory surfaces and experiential memory is something the brain does for us. But it is not just these streams that need to be unified. It's also the many voices of past selves, each with its own constellation of conative and doxastic commitments; these selves are integrated only by self-conscious discipline and work.

I think that this way of understanding things properly captures both the sense in which I—i.e., the subject of *these* self-attributed thoughts and impressions, the thing that thinks when I think and acts when I act—am simple, and the sense in which I am complex. I am complex because I am composed of a collection of subpersonal components, but I speak to the world with a single voice. My voice doesn't decompose, and the attitudes self-attributed in my personal voice belong to me and not my parts.¹⁴ Voices are not made of voice-parts, and the attitudes self-attributed by voices aren't attributed to any part of the system that produces them. One way of putting this is to say that mereology might give the compositional logic of material systems, but it does not give the compositional logic of voices. A community is a collection of people, but a collective voice is not a collection of voices.¹⁵

We treat self-attribution in the Joycean monologue as criterial for a representational state's being properly attributed to the subject, just as we treat reasoning that occurs in the scope of the Joycean monologue as reasoning that is properly attributed to the subject, and being controlled by reasoning that occurs in the Joycean monologue as action that is properly attributed to the subject. All this, together with the fact that the Joycean Machine is the place from which the monologue issues, makes it tempting to identify the self with the Joycean Machine. Now we're in a position to see why that temptation should be resisted. The Joycean Machine is a voice-piece that plays a central unifying role for a collective. But the voice is the voice *of* the collective. To identify the self with the Joycean Machine would be as mistaken as identifying the "we" of the people with the electoral machinery through which it speaks, or the "we" of the Supreme Court justices with its spokesperson.

9. Dynamical Unity

The point of integration of informational streams has to do with dynamics. A voice is a channel for the propagation of information. It can mediate interaction with other systems. In a social setting, a complex system with a voice-piece making public

assertions on its behalf can acquire the normative status of agent, with all of the public commitments and entitlements that entails.¹⁶ This goes for companies and corporations as surely as it does for persons and governments. And when we have a system of collectives communicating through public voices we get an emergent dynamics at the intercollective level: special patterns of interaction that are relatively indifferent to the dynamics at the lower level often not predictable from the laws that govern their components. Again, political communities provide a nice example. Populations band together into national units with governments acting as voice-pieces, giving rise to a dynamics at the international level. Complex feedback and feedforward interactions can make the dynamics at the intercollective level largely autonomous of, and effectively irreducible to, the specific activities of population members. A collective voice can also have an internal role, feeding back into the lower-level organization of the system, guiding the behavior of its components. Think of how laws voted in collectively by a population then have a binding effect on the activity of its members. Or think of a committee that comes together to decide its collective activities and then disbands, leaving each member to carry out its part of the collective plan. The parts of a complex system under the command of a collective voice act with a singleness of intent and purpose that is impossible for the collection of components acting alone. As a general phenomenon, a dynamical link from a higher level to lower levels of organization is the source of most macroscopic order. There are channels for the propagation of information between levels also in self-organizing systems, i.e., systems that don't have voices that self-ascribe intentional attitudes. What is special about systems that have voices in this special sense is that the interaction between the collective and its components is mediated by a self-representational loop, and it is the self-representational loop that creates the space for deliberation, the space within reason operates. The unified voice of the internal monologue does both of these things: it allows us to act as unified agents interacting with other selves, and it also, to the extent that it exercises control over its "constituency," allows them to act as a team, making a coordinated effort in pursuit of a common goal. In sum, the synchronic unity of the thinking subject is the unity of voice and agency wrought by the unifying activity of the Joycean Machine.

Notes

- [1] One of his most important articles on the self is self-consciously titled "The Origins of Selves" (1989) to evoke Darwin's *Origin of the Species*.
- [2] He uses a variety of evocative images for the target, which are not easy to integrate. Aside from the "Cartesian Theater," there is the "Oval Office" in the brain, the "brain pearl," "Central command," and others.
- [3] Much of Dennett's concern centers on the explicit records we keep of our psychological histories. We can recognize that those involve reconstruction and retrospective revision. The purposes for which we keep those records (like, e.g., the records that companies issue to their shareholders explaining corporate actions) may place a low premium on accuracy. Our interest may be less in the reasoning that led to those decisions than on whether we can endorse them and whether they should be repeated.

- [4] The example has also been developed in the literature as an example of distributed cognition by Edwin Hutchins in his wonderful *Cognition in the Wild* (1995).
- [5] A map, for these purposes, is any representation of the landscape rendered in a form that is not relativized to the ship's location.
- [6] This is not to say that it couldn't also represent those goings on: if things broke down, it would need to.
- [7] As we let the time between cycles go to zero, we would have the map simultaneously monitoring and guiding the motions of the ship.
- [8] Computations that transform instrument readings into self-locating coordinates and commands into action are external to the map and not part of the informational stream. They can be added if a need arises to regulate them.
- [9] Dennett (1992) writes, for example:

one of the fundamental tasks performed by the activities of the Joycean Machine is to adjudicate disputes, smooth out transitions between regimes, and prevent untimely *coups d'état* by marshalling the 'right' forces. Simple or overlearned tasks without serious competition can be routinely executed without the enlistment of extra forces, and hence unconsciously, but when a task is difficult... we accomplish it... [with self-manipulations].

These techniques of representing things to ourselves permit us to be self-governors or executives in ways no other creature approaches. We can work out policies well in advance, thanks to our capacity for hypothetical thinking and scenario-spinning; we can stiffen our own resolve to engage in unpleasant or long-term projects by habits of self-reminding, and by rehearsing the expected benefits and costs of the policies we have adopted. (pp. 277–278)

And later: "the broadcasting effect... creates an open forum of sorts, permitting any of the things one has learned to make a contribution to any current problem" (1992, p. 278).

- [10] "The aspirant to a high order of self-control must have the capacity to represent his current beliefs, desires, intentions, and policies in a detached way, as objects for evaluation" (Dennett, 1984, p. 86).
- [11] A map of space reifies structure along the spatial dimension by reproducing the structure of a region of space inside a much smaller part of that region.
- [12] The point is a central theme in the literature on plans and practical reasoning. See, e.g., Bratman (1987).
- [13] Individual members may share the opinions of the group, just as you and I may share opinions, but the opinion of the group is distinct from the opinion of any one or subset of its members.
- [14] This leaves us with a vagueness that I'm happy to acknowledge. What exactly are the contributors to my voice? The informational streams that it unifies originate in the environment, and the distinction between body and environment is soft in ways emphasized in Dennett (1996) and Clark (1997, 2003). We can say the same thing about populations. Who, exactly, is included in the "we" of the people? The informational streams that get resolved into the collective voice are attenuated and there is no obvious terminus. There are the voters, of course, but also the people that voters talk to, and the news media and informal channels through which they get information, and so on. And then there is the question of the collection of parts that are controlled by the government; there are the citizens, of course, the polis, or body politic, but there is a much wider circle of influence affected by more attenuated links and only arbitrary or "legislative" boundaries. All of these sources of vagueness are present, though largely unacknowledged in the case of the self.
- [15] We need to expand our vocabulary for the relations between parts and wholes to make room for collectives in addition to collections, and to interface with our criteria for individuating objects. There is no uniform usage here. Do we want to say you have new 'objects' at the higher level? Or new 'agents' but no new 'objects'? Or is there some other way of describing these cases? These strike me mostly as matters of accounting.

- [16] Of course that status can be revoked: entitlements are hostage to fulfillment of commitments, and nothing has been said about what it takes to acquire the status in the first place.

References

- Bratman, M. E. (1987). *Intention, plans and practical reason*. Cambridge, MA: Harvard University Press.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Clark, A. (2002). That special something: Dennett on the making of minds and selves. In A. Brook & D. Ross (Eds.), *Daniel Dennett* (pp. 187–205). Cambridge, England: Cambridge University Press.
- Clark, A. (2003). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford, England: Oxford University Press.
- Dennett, D. C. (1984). *Elbow Room: The varieties of free will worth wanting*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1989). The origins of selves. *Cogito*, 3, 163–173.
- Dennett, D. C. (1992). *Consciousness explained*. London: Allen Lane.
- Dennett, D. C. (1996). *Kinds of minds*. New York: Basic Books.
- Dennett, D. C. (2003). *Freedom evolves*. London: Allen Lane.
- Dennett, D. C., & Humphrey, N. (1998). Speaking for our selves. In D. C. Dennett (Ed.), *Brainchildren: Essays on designing minds* (pp. 31–58). Cambridge, MA: MIT Press.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Ismael, J. (in press). *The situated self*. Oxford, England: Oxford University Press.