

Daniel Dennett Autobiography Part 2



Daniel C. Dennett reflects on his philosophical life, in this episode from the time he received his first academic post, at the University of California, up to 2003.

University of California Irvine was the perfect setting for me to learn a little philosophy. After all, aside from my undergraduate courses, I had had no serious training in philosophy. (In this regard, if in no other, I am like Quine, who got his PhD in philosophy at Harvard, nominally under Whitehead but in fact, as he once told me, almost entirely self-taught. He began his teaching career after having taken hardly any philosophy courses.) I had never even been a teaching assistant, and now I was going to be an assistant professor, designing and teaching all my own courses, and having teaching assistants of my own – a prospect that both thrilled me and filled me with dread. Would my TAs see through me and denounce me as an impostor? I was Melden's sole appointment that first year of classes at UCI in 1965, and he put me to work teaching the entire undergraduate curriculum, aside from ethics, which he reserved for himself. (The wonderful E. J. Lemmon was appointed during my first year, and came immediately to teach logic, but died of a heart attack before he could take up residence.) I taught Ancient, Medieval, and Modern, epistemology and metaphysics, and, in my second year, a seminar on Anscombe and Taylor on intentionality. I crammed my Aristotle, Augustine, Anselm and Kant, and the others, as well as the more accessible secondary sources, and stayed a day or two ahead of my eager but unsophisticated students. I blush to consider what I must have imparted, but hopefully my genuine excitement about these brand new texts made up for the naïvete and just plain ignorance with which I addressed them.

It was time to turn my D.Phil dissertation into articles and a book. I revised the first chapter and sent it out as a journal article. A dozen submissions: a dozen rejections, with many revisions in between. Then Wilfrid Sellars, editor of *Philosophical Topics*, wrote me a nice letter saying that he was intrigued by the draft I had sent. He thought it would be fine once I clarified a few foggy points. I sent him a clarified version within the week, and he wrote back to say that now it was clear what I was doing, he thought it wasn't publishable! A few more rejections and I gave up on that chapter and started several other projects, with no greater success. Perhaps I wasn't going to make it as a philosopher after all.

One day Julian Feldman, an Artificial Intelligence researcher at UCI, came storming into my office with a copy of Hubert Dreyfus's notorious memo for think-tank RAND, 'Alchemy and Artificial Intelligence'. What did I make of it? I read it and said I disagreed quite fundamentally with it. "Write up your rebuttal, please, and get it published!" Why not? I wrote 'Machine Traces and Protocol Statements' and promptly published it in *Behavioral Science* (1968). My first publication – and my career as philosopher-laureate of AI had begun. I'd

already been attracted to the field by Alan Ross Anderson's pioneering anthology, *Minds and Machines*, and found at Irvine a small group of AI researchers who invited me to join them. Allen Newell came through town to give some talks. He struck up a lively conversation with me, and I was hooked. Other colleagues at Irvine, in particular the psychobiologist James McGaugh, struck by my knowledge of and interest in theories of learning in neural systems, also took a vigorous interest in further educating me and getting me thinking about their work and its problems. Nor were my philosophical colleagues shy about contributing to my education. Joe (Karel) Lambert had been an experimental psychologist before he turned to logic, and he was a knowledgeable and resourceful defender of Skinnerian behaviorism, quick to puncture my superficial knowledge and rub my nose in the experimental literature, not just the manifestos of Skinner and his fellow ideologues.

In the summer of 1967 I sent a revised version of my dissertation to the famous Routledge & Kegan Paul series, *The International Library of Philosophy and Scientific Method*. This series of books, with their red covers and yellow dust jackets, included most of my favorite books in philosophy: Wittgenstein's *Tractatus*, Smart's *Philosophy and Scientific Realism* and Sellars' *Science, Perception and Reality*, for instance. A year passed without a word from the new editor, Ted Honderich, who had taken over from A.J. Ayer. I didn't dare upset the applecart by complaining about his unresponsiveness. Finally, when I knew I was off to Oxford for a quarter sabbatical in the fall of 1968, I wrote a timid inquiry to Honderich, who discovered that the manuscript had been mislaid by the referee to whom he had sent it. Honderich retrieved it, read it himself and forthwith accepted it, pending revisions which I hastened to complete that autumn in Oxford. I was in heaven. But still, I couldn't talk about it to other philosophers.

My problem was that my way of approaching the then-standard issues in the philosophy of mind was too eccentric, too novel, to afford easy entry into a discussion. When somebody asks you what you're working on, you usually can't back them into a corner and harangue them for a couple of hours about your project, and I could imagine no more modest framing job. After all, my attempts to publish the first chapter of my dissertation showed that the first ideas I needed to get across were bound to be misunderstood – had already been misunderstood in half a dozen versions by some of the best philosophers of mind in the field. So I was a very lonely and uncertain philosopher those first few years at Irvine, spending more happy hours talking AI or psychobiology than philosophy. In spite of my presumably sterling pedigree as a student of Quine and Ryle I felt like an outsider, a dark horse one should probably not bet

on. The acceptance of the manuscript by Honderich, and his further invitation to write an essay on free will ('Mechanism and Responsibility', in which I introduced the terminology of the intentional stance and intentional systems) gave me new confidence, however. I began to get assertive with my very vigorous and somewhat combative colleagues in the Philosophy Department. Joe Lambert, Gordon Brittan, Stan Munsat, Jack Vickers and I

participated in a seminar on Hintikka's *Knowledge and Belief*, in which I trotted out the intentional

stance in its first version. They all jumped on me, except for Vickers, who saw what I was up to right away, and encouraged me to persist. (The intentional stance is the strategy of interpreting an object as an agent with beliefs, desires and rationality. You can adopt the intentional stance towards a person, an animal, a thermostat – which 'wants' to maintain a certain temperature, and regularly updates its 'belief' about what the current temperature is – or, more interestingly, a chess-playing computer, which 'knows' the rules, has true 'beliefs' about the positions of the pieces on the board and 'wants' to win.) A later joint seminar, reading Jerry Fodor's *Psychological Explanation*, was even more raucous. I remember that we terrified the graduate students with our unrelenting attacks on each other, pouncing with glee on presumed inconsistencies and stupidities, descending even to taunts on occasion. But in fact we had great respect and affection for each other, so while we set a bad example, we did no harm to each other, and managed, in fact, to anneal some pretty good philosophy in our white hot ovens.

When my book *Content and Consciousness* came out in 1969, UCI put me up for early tenure, and asked whom they should get to referee the book. Quine, I said, since I had to know whether he liked or hated it, and here was my one sure chance to get him to read it. He was the first reader of the book aside from its author and editor, and he liked it, aside from the embarrassing factual mistakes ('Giorgone' means Big George, not Little George; 'twas Ponce de Leon, not de Soto, who searched for the Fountain of Youth!) From the vantage point of more than a third of a century later, my opinion is that my curious methods paid off handsomely in that first book. I secured a breakthrough position that has proved remarkably stable and fecund over the years. I've been happily turning the crank ever since, making minor adjustments to the defenses and generating more and more offspring ideas, all showing their ancestry: I'm what you get when you cross Quine with Ryle and add some cognitive science.

In six years at UCI, I never gave a talk anywhere in California. My first invitation came from Princeton in December of 1970, and I presented 'Intentional Systems' to a daunting audience, including Alonzo Church, Donald Davidson, Dick Rorty, David Lewis, Tom Nagel, and all the others. They liked it; the discussion was fine, and I decided I had to move east, where it seemed to me the action was. Susan and I had both summered in Maine as children, and we bought a farm in Blue Hill, Maine,

while visiting our folks in the Boston area at Christmas. Tufts University [in Massachusetts] beckoned a few months later, in 1971, and I answered. I've never regretted it. For thirty years and more, Hugo Bedau and the happy crew he assembled at Tufts have been a high-quality, fair, faction-free, supportive, decent department. Judging by the perennial grumbles and moans of my friends in strife-ridden philosophy departments

around the world, I reckon that I have been spared thousands of prime-time hours of distracting rumination and counterplotting – enough time to write several books. By teach-

ing in a department that didn't have a PhD program, I was also spared the anxious drudgery of carrying mediocre graduate students across the finish line and into tenure track jobs. Even in the top PhD schools there are quite a few problematic dissertationers, and I surely would have had my share to shepherd. Since 1986, I have been in the enviable position of having, in the Center for Cognitive Studies, a funded dissertation fellowship or post-doc to award, but only when I wished, and the alumni of this position have amply borne out my judg-

“ I find that the cutting-edge ideas of science are typically more beautiful, more delightful to contemplate, than most of the ideas I encounter in philosophy journals ”

You Are Here Madden.




ment. The Center for Cognitive Studies was Tufts' response to my one serious thought of moving, when Wilfrid Sellars retired and Pittsburgh offered me his chair. In addition to the annual fellowship for a collaborator/student of my choice, Tufts gave me a reduced teaching load, clerical assistance, and an administrative budget that relieved the burden I was putting on the regular department budget. This has provided me with more than my fair share of time – the one commodity in limited supply no matter how rich you are.

During the 70s I learned more AI and psychology, and wrote a variety of papers that I decided would have a greater impact if gathered into a collection. Just as I was about to go scouting for a publisher, I was paid a visit by an elegant and charming couple, Harry and Betty Stanton, founders of a new publishing venture, Bradford Books. Harry was a veteran editor with experience in large publishing houses who was eager to publish high-quality books on his own – from their house in Vermont ski country! He had hit upon the general area of what would soon be called cognitive science, and on a visit to MIT he'd been told by Noam Chomsky and Ned Block that I would be someone to talk to. I liked the Stantons' pitch and handed them the stack of papers I'd put together to make my collection, telling them that I would publish with them if they would bring the book out simultaneously in cloth and paperback and let me have veto power over the paperback price. I disapproved of the outrageous prices university presses were asking for their books, and wanted this book to be available immediately at a reasonable price to students. So sure was I that this was a good idea that I was prepared to pay for it by foregoing royalties on the first 3000 copies – more than most philosophy books ever sell. They mulled and conferred with the Kitchers, then at University of Vermont (the nearest good philosophers to their ski house) and agreed.

Since Harry wanted to start his backcountry publishing venture with a splash, he devised a clever advertising campaign, with Edward Gorey artwork and amusing fliers; and when *Brainstorms* came out in 1978, it attracted attention across all the fields that composed cognitive science. Soon the Stantons' main problem was dealing with the deluge of manuscripts from would-be Bradford Books authors. But then Harry got cancer, and while he soon had it under control, he had to abandon his plans of keeping Bradford Books a stand-alone company. MIT Press bought the imprint and hired Harry and Betty as its editors, and they proceeded to dominate cognitive science publishing for decades, especially its philosophical side.

One of the side benefits for me was that as I became better known among researchers in these fields, I found that they were willing, even eager, to become my informants. Over the years I have compiled a priceless tutorial education from the leaders in linguistics, AI, psychology, neuroscience and evolutionary biology. They typically expected a philosopher to be complacently ignorant of their fields, so what little I did know usually pleasantly surprised them, and encouraged them to teach me more. What I was early to recognize was that scientists, however brusque and dismissive of philosophy they first appear, typically aspire to answer the Big Questions, or at least contribute to the answers. If a diplomatic and open-minded philosopher will learn something of their field and take them






seriously, ignoring the trivialities – the misunderstanding of philosophical jargon, the naïvetés of formulation, and the awkward moves – they will provide a bounty of philosophical insights that only need a little translating and adjusting. Besides – and I might as well put this bluntly – I find that the



cutting-edge ideas of science are typically more beautiful, more delightful to contemplate, than most of the ideas I encounter in philosophy journals. We philosophers tend to wander back and forth between cramped, blinkered, nibbling exercises on the one hand, and advertisements for grandiose but half-baked visions on the other. The fact that the axes that scientists grind can be shown to be right or wrong, often quite quickly, gives them a heft and sharpness that philosophical axes seldom achieve. Thus my vision of the philosopher's role has been gradually reshaped, ever more in the direction of providing the conceptual clarifications and underpinnings for theories that are testable, empirical, scientific. This was also Quine's view, of course; but he didn't get much chance to actually do philosophy in this vein.

I spent the year 1979-80 at the Center for Advanced Study in the Behavioral Sciences in Palo Alto, as part of a six-person group working with John McCarthy on AI and philosophy. In spite of the fact that one could hardly pick four AI researchers more attuned to philosophy than those present (McCarthy, Patrick Hayes, Zenon Pylyshyn and Robert Moore), or two philosophers more attuned to AI than John Haugeland and I, we often talked past each other without realizing it, and it was late in the year when I finally figured out some of what was driving the AI gang. Interdisciplinary communication has many pitfalls, most of which I have fallen in and subsequently learned to avoid. I often cringe when I witness philosophers boldly charging into the interdisciplinary fray, uttering words that work just fine in philosophy seminars but are guaranteed to irritate, mislead and baffle their non-philosophical audience.

I also worked hard that year on a long paper I'd begun the year before in Bristol, where Steve Stich and I had been Fulbright Fellows in a philosophy and psychology group put together by Andrew Woodfield. Stich and I had hoped to co-



author a paper on the troublesome topic of *de re* and *de dicto*, wide and narrow content, and related issues about propositional attitudes. We were both deeply skeptical of the received wisdom, but in the end we couldn't agree on how to launch an alternative trajectory, and went our separate but largely parallel ways. There were just too many loose ends and dubious presumptions keeping the whole enterprise afloat, and I found myself making and breaking, building and rebuilding, reaching back further and further to find some ground that wasn't too slippery to stand on. This was very, very hard work for me – the hardest work I have ever done, I think. In the end, I settled for a version that was published as 'Beyond Belief', which raised some challenges for the propositional attitude folks that have never been properly answered, I think. I now can see better ways of dealing with the issues, growing in part out of Ruth Millikan's work, but working this out is not a high priority for me. Clearing up the confusions of philosophers seems less important to me than clearing up the philosophical confusions of people working in cognitive science. (I'm conferring with roboticists these days about the task of giving robots concepts of the things in their worlds which will permit them to recognize them and deal with them intelligently – the real phenomenon which philosophers misname *de re* belief and confuse with various fantasies. It will be interesting to see if we can make progress in this end-run around the philosophical swamps.)

During the year in Palo Alto I was visited by Douglas Hofstadter, who had just published *Gödel, Escher, Bach* as well as a laudatory review of *Brainstorms* in the *New York Review of Books*. He introduced himself and promptly proposed that we collaborate on an anthology of think-pieces about the mind, drawn from all sources – science fiction, philosophy, science, whatever. I was initially reluctant, but he was persistent – thank goodness – and the resulting volume, *The Mind's I* (1981), was a huge success. It changed my life.

Hofstadter is one of the deepest, most inventive thinkers I have ever encountered. He loves the big philosophical

questions, but is generally impatient with the way philosophers tackle them. His dim view of philosophical work is occasionally just wrong, in my opinion, but not always; and I found that I sometimes didn't really believe my half-hearted defenses of some of our traditions. The Mount Everest reason for tackling some huge pile of philosophical confusion ('because it's there') appealed less and less to me, and alternative paths around what I was sure were philosophical dead-ends became ever more attractive. I saw clearly enough that if I turned my back on some of the hot philosophical issues, those who were engaged in them would quite appropriately turn their back on me, but it was a fate I was prepared to endure. With any luck, I might be able to oblige them to play my games as well as their own.


One of the authors Hofstadter introduced to me was Richard Dawkins, and I soon became addicted to the kind of thinking that goes on in contemporary evolutionary theory. Such tantalizing models, such a beautiful set of moves, which could take you, in crystal-clear steps, all the way from molecules to Meaning! I had long been an ardent if entirely amateur Darwinian. Now I began to educate myself in the finer points, consuming the works of George Williams, John Maynard Smith, William Hamilton, Robert Trivers and others. One of the first things that struck me – as it has struck others



when they get into this literature – was that my previous chief informant on matters Darwinian, Stephen Jay Gould (a distant cousin of Doug Hofstadter, by the way, who introduced us), seriously misrepresented the evolutionary theory he disagreed with in his work. I attended his seminar on evolution and cognition at Harvard, and found it so troubling that in

private conversation I tried to persuade him to give a more evenhanded presentation of the views he disliked so much, but to no avail. I couldn't bear to continue in the seminar, but we were still on friendly terms, and he and Dawkins both made guest appearances in my Tufts seminar on philosophy and evolution. Gould, however, was not used to being questioned the way my Tufts students questioned him. He took offense and cut short the class. My students were appalled, and wanted to know if they had done wrong. No, I told them; I had never been prouder of Tufts students in my life. After that, although Gould and I were still cordial – for instance, in the memorable television round-table made in the Netherlands, *A Glorious Accident* – the wall descended between us.

From the mid 1980s, my attention was divided between evolutionary biology on the one hand and cognitive science on the other. Since the work being done in both fields struck me as both more fascinating and more fruitful – more philosophically enlightening – than almost any work I saw being done in philosophy, I found myself letting go of the philosophical grapevines and collaborating with people in other fields. I had begun working with the neuropsychologist Marcel Kinsbourne, joining him on rounds at various hospitals to see first hand some of the pathologies I was trying to account for in my



Stephen Jay Gould with DCD

revised theory of consciousness. When the psychologist Nick Humphrey came to work with me, I introduced him to the linguist Ray Jackendoff, and the four of us began meeting regularly to discuss our various consciousness projects. Coming from four different directions, with different methods, styles and aspirations, we nevertheless made a good team.

I could never have written *Consciousness Explained* (1991) without their help, just as I could never have written *Darwin's Dangerous Idea* (1995) without the help of such evolutionists as Richard Dawkins, David Haig, Ernst Mayr, E.O. Wilson and Steve Pinker. Both books are, in one sense, books of philosophy; but both also aspire to contribute more directly to science by defending specific empirical theories, of consciousness and of evolution, and more of my energies in the last decade have been devoted to pursuing these theoretical developments with people in those fields than to working with philosophers.

In both books, by the way, I was pursuing a strategy I had been discussing for years with friends such as Stevan Harnad, founder and editor of *Behavioral and Brain Sciences*, as accomplished an interdisciplinarian as you can find. In any field, when experts talk to experts, the worst *faux pas* is over-explaining, which is insulting – so experts routinely err on the side of under-explaining. The result is that while experts may impress each other, they tend to talk past each other, especially in interdisciplinary contexts. The cure, I surmised, was to have experts address their talk to novices while being ‘overheard’ by their fellow experts. This way the elementary level of explanation has a suitable excuse, while the onlooking experts can, without embarrassment, benefit from a discussion that is pitched at what is in fact just the right level. My trade books are officially written for the proverbial educated lay audience, and many of these readers find that, with a struggle, they can follow the arguments just fine. But they are not my real target. I’m writing for my fellow experts, in philosophy, cognitive science, evolutionary biology and other fields, in terms they can readily understand. I think only some philosophers among them turn up their noses at this ploy. These philosophers are still under the impression that if you don’t have to struggle to read a book, it can’t be worth reading. I suspect it’s their way of proving to themselves that they’re trained professionals: there is something they can do that untrained people just can’t: understand a philosophy book.

When I first read Descartes, in my freshman year in college, it struck me hard that his idea of the self as a *res cogitans* – the immaterial source of intentional action and the locus of all mattering – just had to be a mistake, and a huge one. The idea that free will and responsibility could hinge on such a metaphysical singularity was, I thought, preposterous. So from the outset I was a compatibilist on free will, convinced that one way or another it should be possible to reconcile at least a generous portion of the tradition about free will – the free will worth wanting – with the manifest fact that we are, as I now like to put it, composed of trillions of mindless robots and nothing else. When Oxford invited me to give the John Locke Lectures in 1983, I found myself in pretty good position to pull my various reflections on this topic together. The result was *Elbow Room* (1984).

I had attended Locke Lectures in Oxford on several occasions, both as a graduate student and as a visitor on sabbatical, and I knew the tradition: the American professor giving the talks, intent on demonstrating to Oxford a high degree of technical expertise and advanced philosophical knowledge, reads a series of lectures bristling with difficulties that only the most intrepid and well-versed experts can follow. The audience does a Zeno, halving its population each week – or worse – so that by the time lecture six rolls around, a few graduate students must be dragooned to sit in the nearly empty hall so as not to embarrass the speaker. I decided to be different, going out of my way to make my lectures accessible even to non-philosophers who might attend. The result was standing room only for all six lectures.

I paid a price, of course. One eminent Oxford philosopher was overheard saying as he left he was damned if he could learn anything from somebody who could attract such hordes week after week. That attitude has intensified and spread in some philosophical circles since then, as I have steadily enlarged my non-philosophical audience, thereby persuading more than a few philosophers that I cannot be doing anything worth considering. I found, for instance, that I will have to pull the serious and original arguments out of *Darwin's Dangerous Idea* and *Kinds of Minds* and restate them in journal articles to get expert philosophers to pay attention to them.

It's not a bad arrangement: while the experts grind away at their for-philosophers-only cottage industries, their undergraduate students gratefully read my books and articles and absorb a great deal which stands by them when they enter the dark satanic mills of graduate school. Concepts such as the personal/subpersonal level distinction, folk psychology, and intuition pump get put to good use without attribution, along with the more noticeably Dennettian concepts of the intentional stance and design stance, the self as a center of narrative gravity, and the Cartesian theater, among others. Ryle was right: head-to-head argumentative combat – definition, refutation and rebuttal – is not the only way to influence philosophical thinking, or even the best.

© PROF. DANIEL C. DENNETT 2008

Dan Dennett is Co-Director of the Center for Cognitive Studies and is Austin B. Fletcher Professor of Philosophy at Tufts University. His latest book is *Breaking the Spell* (Viking, 2006).

• This article was written in 2003. In our next issue you can read a Postscript by Professor Dennett bringing his story up to date.